

The importance sampling technique for understanding rare events in Erdős-Rényi random graphs

Shankar Bhamidi* Jan Hannig† Chia Ying Lee‡ James Nolen§

Abstract

In dense Erdős-Rényi random graphs, we are interested in the events where large numbers of a given subgraph occur. The mean behavior of subgraph counts is known, and only recently were the related large deviations results discovered. Consequently, it is natural to ask, can one develop efficient numerical schemes to estimate the probability of an Erdős-Rényi graph containing an excessively large number of a fixed given subgraph? Using the large deviation principle we study an importance sampling scheme as a method to numerically compute the small probabilities of large triangle counts occurring within Erdős-Rényi graphs. We show that the exponential tilt suggested directly by the large deviation principle does not always yield an optimal scheme. The exponential tilt used in the importance sampling scheme comes from a generalized class of exponential random graphs. Asymptotic optimality, a measure of the efficiency of the importance sampling scheme, is achieved by a special choice of the parameters in the exponential random graph that makes it indistinguishable from an Erdős-Rényi graph conditioned to have many triangles in the large network limit. We show how this choice can be made for the conditioned Erdős-Rényi graphs both in the replica symmetric phase as well as in parts of the replica breaking phase to yield asymptotically optimal numerical schemes to estimate this rare event probability.

Keywords: Erdős-Rényi random graphs; exponential random graphs; rare event simulation; large deviations; graph limits.

AMS MSC 2010: 65C05; 05C80; 60F10.

Submitted to EJP on March 24, 2013, final version accepted on September 30, 2015.

1 Introduction

In this paper we study the use of importance sampling schemes to numerically estimate the probability that an Erdős-Rényi random graph contains an unusually large number of triangles. A simple graph X on n vertices can be represented as an element

*Department of Statistics and Operations Research, University of North Carolina, USA.

E-mail: bhamidi@email.unc.edu

†Department of Statistics and Operations Research, University of North Carolina, USA.

E-mail: hannig@email.unc.edu

‡Department of Statistics and Operations Research, University of North Carolina, USA.

E-mail: Chia_Ying_Lee@alumni.brown.edu

§Mathematics Department, Duke University, USA.

E-mail: nolen@math.duke.edu

of the space $\Omega_n = \{0, 1\}^{\binom{n}{2}}$. A graph $X \in \Omega_n$ will be denoted by $X = (X_{ij})_{1 \leq i < j \leq n}$ with the entry X_{ij} indicating the presence or absence of an edge between vertices i and j . For a given edge probability $p \in [0, 1]$, an Erdős-Rényi random graph $\mathcal{G}_{n,p}$ is a graph on n vertices such that any edge is independently connected with probability p . We shall use $\mathbb{P}_{n,p}$ to represent the probability measure on Ω_n induced by the Erdős-Rényi graph $\mathcal{G}_{n,p}$. The probability of a fixed graph X under the measure $\mathbb{P}_{n,p}$ can be explicitly computed as

$$\mathbb{P}(G_{n,p} = X) = \mathbb{P}_{n,p}(X) = \prod_{i < j} p^{X_{ij}} (1-p)^{1-X_{ij}} = (1-p)^{\binom{n}{2}} e^{h_p E(X)} \quad (1.1)$$

where

$$h_p := \log \frac{p}{1-p} \quad (1.2)$$

and $E(X) := \sum_{i < j} X_{ij}$ is the number of edges in X . Let $T(X)$ denote the number of triangles in graph X :

$$T(X) = \sum_{1 \leq i < j < k \leq n} X_{ij} X_{jk} X_{ik}.$$

Also, let the event

$$W_{n,t} = \left\{ X \in \Omega_n \mid T(X) \geq \binom{n}{3} t^3 \right\}$$

denote the upper tails of triangle counts. Consider an Erdős-Rényi random graph $\mathcal{G}_{n,p}$. For p fixed, one can show that $\mathbb{E}[T(\mathcal{G}_{n,p})] \sim \binom{n}{3} p^3$ as $n \rightarrow \infty$. For $t > p$, the main aim of this paper is the following question: can one develop efficient numerical schemes to estimate the probability

$$\mu_n = \mathbb{P} \left(T(\mathcal{G}_{n,p}) \geq \binom{n}{3} t^3 \right) \quad (1.3)$$

that $\mathcal{G}_{n,p}$ has an atypically large number of triangles? Before addressing such questions, one first needs to understand the structure of such random graphs, conditioned on this rare event, more precisely the large deviation rate function for such events. The last few years have witnessed a number of deep results in understanding such questions including upper tails of triangle counts, along with more general subgraph densities (see e.g., [3, 5, 6, 8, 9, 13, 17]). In the dense graph case, where the edge probability p stays fixed as $n \rightarrow \infty$, [6] derived a large deviation principle (LDP) for the rare event $\{T(\mathcal{G}_{n,p}) \geq \binom{n}{3} t^3\}$, showing that for t within a certain subset of $(p, 1]$,

$$\mathbb{P} \left(T(\mathcal{G}_{n,p}) \geq \binom{n}{3} t^3 \right) = \exp \left(-n^2 I_p(t) (1 + O(n^{-1/2})) \right) \quad (1.4)$$

where the rate function $I_p(t)$ is given by

$$I_p(t) = \frac{1}{2} \left(t \log \frac{t}{p} + (1-t) \log \frac{1-t}{1-p} \right). \quad (1.5)$$

More recently [8] showed a general large deviation principle for dense Erdős-Rényi graphs, using the theory of limits of dense random graph sequences developed recently by Lovasz et al. [14, 16, 15, 3]. When specialized to upper tails of triangle counts, they show that there exists a rate function $\phi(p, t)$

$$\frac{1}{n^2} \log \mathbb{P} \left(T(\mathcal{G}_{n,p}) \geq \binom{n}{3} t^3 \right) \rightarrow -\phi(p, t), \quad \text{as } n \rightarrow \infty. \quad (1.6)$$

The function $\phi(p, t)$ coincides with $I_p(t)$ for a certain parameter range of (p, t) , and is described in more detail in (2.5). The exponential decay of the probability of the event

of interest makes it difficult to estimate this probability even for moderately large n . Direct Monte Carlo sampling is obviously intractable. The central strategy of importance sampling is to sample from a different probability measure, the *tilted measure*, under which the event of interest is no longer rare; one obtains more successful samples falling in the event of interest but each sample must then be weighted appropriately according to the Radon-Nikodym derivative of the original measure against the tilted measure. Importance sampling techniques have been used in many other stochastic systems, such as SDEs and Markov processes and queuing systems, see e.g [4, 21, 12, 10, 2] and the references therein. In particular, when a large deviations principle is known for the stochastic system, the tilted measure commonly used is a change of measure arising from the LDP. However, not every tilted measure associated with the LDP works well. It is well known that a poorly chosen tilted measure can lead to an estimator that performs worse than Monte Carlo sampling, or whose variance blows up [11]. Thus, a careful choice of tilted measure is of utmost importance. Before describing the relevant tilts we formally define our aims.

1.1 Importance sampling and asymptotic optimality

If $\{X^k\}_{k=1}^\infty \subset \Omega_n$ is a sequence of Erdős-Rényi random graphs generated independently from $\mathbb{P}_{n,p}$, then for any integer $K \geq 1$,

$$M_K = \frac{1}{K} \sum_{k=1}^K \mathbf{1}_{W_{n,t}}(X^k)$$

is an unbiased estimate of μ_n . By the law of large numbers, $M_K \rightarrow \mu_n$ with probability one as $K \rightarrow \infty$. Although this estimate of μ_n is very simple, the relative error is

$$\frac{\sqrt{\text{Var}(M_K)}}{\mathbb{E}(M_K)} = \frac{\sqrt{\mu_n - (\mu_n)^2}}{\mu_n \sqrt{K}},$$

which scales like $(K\mu_n)^{-1/2}$ as $\mu_n \rightarrow 0$. Hence the relative error may be very large in the large deviation regime where $\mu_n \ll 1$, unless we have at least $K \sim O(\mu_n^{-1})$ samples. Therefore, it is desirable to devise an estimate of μ_n which, compared to this simple Monte Carlo estimate, attains the same accuracy with fewer number of samples or lower computational cost.

Importance sampling is a Monte Carlo algorithm based on a change of measure. Suppose that $\mathbb{P}_{n,p}$ is absolutely continuous with respect to another measure \mathbb{Q} on Ω_n with

$$\frac{d\mathbb{P}_{n,p}}{d\mathbb{Q}} = Y^{-1} : \Omega_n \rightarrow \mathbb{R}.$$

Then we have

$$\mu_n = \mathbb{E}[M_K] = \mathbb{E} \left[\frac{1}{K} \sum_{k=1}^K \mathbf{1}_{W_{n,t}}(X^k) \right] = \mathbb{E}_{\mathbb{Q}} \left[\frac{1}{K} \sum_{k=1}^K \mathbf{1}_{W_{n,t}}(\tilde{X}^k) Y^{-1}(\tilde{X}^k) \right] \tag{1.7}$$

where $\mathbb{E}_{\mathbb{Q}}$ denotes expectation with respect to \mathbb{Q} , and we now use $\{\tilde{X}^k\}_{k=1}^\infty$ to denote a set of random graphs sampled independently from the new measure \mathbb{Q} . If we define

$$\tilde{M}_K = \frac{1}{K} \sum_{k=1}^K \mathbf{1}_{W_{n,t}}(\tilde{X}^k) Y^{-1}(\tilde{X}^k), \tag{1.8}$$

then \tilde{M}_K is also an unbiased estimate of μ_n , and the relative error is now:

$$\frac{\sqrt{\text{Var}_{\mathbb{Q}}(\tilde{M}_K)}}{\mathbb{E}_{\mathbb{Q}}(\tilde{M}_K)} = \frac{\sqrt{\mathbb{E}_{\mathbb{Q}}[(\mathbf{1}_{W_{n,t}}(\tilde{X}) Y^{-1}(\tilde{X}))^2] - (\mu_n)^2}}{\mu_n \sqrt{K}}, \tag{1.9}$$

Formally this is optimized by the choice $Y = (\mu_n)^{-1} \mathbf{1}_{W_{n,t}}(X)$, in which case the relative error is zero. Such a choice for Q is not feasible, however, since normalizing Y would require a priori knowledge of $\mu_n = \mathbb{P}_{n,p}(W_{n,t})$. Intuitively, we should choose the tilted measure Q so that $\tilde{X}_k \in W_{n,t}$ occurs with high probability under Q .

We will refer to Y^{-1} as the importance sampling weights, and Q as the *tilted measure*, or tilt. If Q arises naturally as the measure induced by a random graph \mathcal{G}_n , we will also refer to \mathcal{G}_n as the tilt. In the cases where a large deviation principle holds, it gives us an estimate of the relative error in the estimate \tilde{M}_K . For any fixed K , it is clear from (1.9) that minimizing the relative error is equivalent to minimizing the second moment $\mathbb{E}_{Q_n}[(\mathbf{1}_{W_{n,t}} Y^{-1})^2]$. Since Jensen's inequality implies that

$$\mathbb{E}_{Q_n}[(\mathbf{1}_{W_{n,t}} Y^{-1})^2] \geq \mathbb{E}_{Q_n}[\mathbf{1}_{W_{n,t}} Y^{-1}]^2,$$

we have the following asymptotic lower bound:

$$\liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{Q_n}[(\mathbf{1}_{W_{n,t}} Y^{-1})^2] \geq -2\phi(p, t). \tag{1.10}$$

Thus, the presence of a large deviation principle for the random graphs $\mathcal{G}_{n,p}$ as $n \rightarrow \infty$, leads to a way to quantify the efficiency of the importance scheme in an asymptotic sense, as is done in other contexts [4].

Definition 1.1. A family of tilted measures Q_n on \mathcal{W} is said to be asymptotically optimal if

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{Q_n}[(\mathbf{1}_{W_{n,t}} Y^{-1})^2] = -2 \inf_{f \in W_{n,t}} [I_p(f)].$$

In contrast, the second moment of each term in the simple Monte Carlo method satisfies

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{P_n}[\mathbf{1}_{W_{n,t}}^2] = -\phi(p, t) > -2\phi(p, t).$$

Thus, the simple Monte Carlo method is not asymptotically optimal. Observe that Jensen's inequality for conditional expectation implies

$$\begin{aligned} Q_n(W_{n,t})^{-1} &= P_n(W_{n,t})^{-1} \left(\frac{\mathbb{E}_{P_n}(\mathbf{1}_{W_{n,t}} Y)}{P_n(W_{n,t})} \right)^{-1} \\ &\leq P_n(W_{n,t})^{-2} \mathbb{E}_{P_n}(\mathbf{1}_{W_{n,t}} Y^{-1}) = P_n(W_{n,t})^{-2} \mathbb{E}_{Q_n}(\mathbf{1}_{W_{n,t}} Y^{-2}). \end{aligned} \tag{1.11}$$

So, if Q_n is asymptotically optimal, we must have

$$\liminf_{n \rightarrow \infty} \frac{1}{n^2} \log Q_n(W_{n,t}) \geq \liminf_{n \rightarrow \infty} \frac{2}{n^2} \log P_n(W_{n,t}) + \liminf_{n \rightarrow \infty} \frac{-1}{n^2} \log \mathbb{E}_{Q_n}(\mathbf{1}_{W_{n,t}} Y^{-2}) = \mathbb{0} \tag{1.12}$$

which is consistent with the intuition that a good choice of Q_n should put $\tilde{X}_k \in W_{n,t}$ with high probability.

To understand in this context the tilts that could be relevant, let us now describe in a little more detail, properties of the rate function 1.6 as well as structural results of the Erdős-Rényi model conditioned on rare events and their connections to a sub-family of the famous exponential random graph models.

1.2 Edge and triangle tilts

In this article we consider tilted measures within a family of exponential random graphs $\mathcal{G}_n^{h,\beta,\alpha}$. For parameters $h \in \mathbb{R}$, $\beta \geq 0$, and $\alpha > 0$, these exponential random graphs are defined via the Gibbs measure, $Q_n = Q_n^{h,\beta,\alpha}$ on the space of simple graphs on n vertices, where

$$Q_n^{h,\beta,\alpha}(X) \propto e^{H(X)}, \quad \text{where } H(X) = hE(X) + \frac{\beta}{n} \left(\frac{n^3}{6} \right)^{1-\alpha} T(X)^\alpha. \tag{1.13}$$

$E(X)$ is the number of edges in graph X . If $\beta = 0$, and $h = h_q = \log \frac{q}{1-q}$ for some $q \in (0, 1)$, then $\mathcal{G}_n^{h_q, 0, \alpha}$ is an Erdős-Rényi graph with edge probability q (notice that α is irrelevant when $\beta = 0$). In particular, the original graph $\mathcal{G}_{n,p}$ is an exponential random graph with parameters $h = h_p$ and $\beta = 0$.

Given the rare event problem, $\mathcal{G}_{n,p}$ conditioned on $T(\mathcal{G}_{n,p}) \geq \binom{n}{3}t^3$, which we shall henceforth parameterize by (p, t) , we will focus on two strategies for choosing the tilted measure. The first is to set $\beta = 0$ and $h = h_q$ for some $q > p$. The resulting tilted measure $\mathbb{Q}_n^{h_q, 0, \alpha}$ will be called an *edge tilt*; compared to the original measure for $\mathcal{G}_{n,p}$, this tilt simply puts more weight on edges. The second strategy is to set $h = h_p$ but vary $\beta > 0$ and $\alpha > 0$. We refer to the resulting tilted measure $\mathbb{Q}_n^{h_p, \beta, \alpha}$ as a *triangle tilt*; compared to the original measure, this tilt puts more weight on triangles, while leaving $h = h_p$ unchanged.

That the two tilts above are natural candidates for the importance sampling scheme, can be reasoned in light of the following concept of when two graphs are alike. In [8] it is shown that for the range of (p, t) where one has (1.4), the Erdős-Rényi graph $\mathcal{G}_{n,p}$ conditioned on the rare event $\{T(\mathcal{G}_{n,p}) \geq \binom{n}{3}t^3\}$ is *asymptotically indistinguishable* from another Erdős-Rényi graph $\mathcal{G}_{n,t}$ with edge probability t , in the sense that the typical graphs in the conditioned Erdős-Rényi graph resembles a typical graph drawn from $\mathcal{G}_{n,t}$ when n is large. (Asymptotic indistinguishability is explained more precisely at (2.7).) Thus, choosing the tilted measure to resemble the typical conditioned graph is a natural choice. While it may seem plausible for any $t > p$ that the conditioned graph resembles another Erdős-Rényi graph, since $\mathbb{E}[T(\mathcal{G}_{n,t})] \sim \binom{n}{3}t^3$ as $n \rightarrow \infty$, it is not always the case. Depending on p and t , it may be that the graph $\mathcal{G}_{n,p}$ conditioned on the event $\{T(\mathcal{G}_{n,p}) \geq \binom{n}{3}t^3\}$ tends to form cliques and hence does not resemble an Erdős-Rényi graph. When the conditioned graph does resemble an Erdős-Rényi graph, we say that (p, t) is in the *replica symmetric* phase. On the other hand, when the conditioned graph is not asymptotically indistinguishable from an Erdős-Rényi graph we say that (p, t) is in the *replica breaking* phase. (See Definition 2.2.)

The main question we wish to address is: given the parameters (p, t) for the rare event problem, how can we choose the tilt parameters (h_q for the edge tilt, or β and α for the triangle tilt) so that the resulting importance sampling scheme is asymptotically optimal? And, can an optimal importance sampling scheme be constructed for all values of (p, t) ?

Regarding the edge tilt, our first result (Prop 3.4) is that the edge tilt $\mathbb{Q}_n^{h_q, 0, \alpha}$ can be asymptotically optimal only if $h_q = h_t$ (i.e. $q = t$). This is not very surprising since $\mathbb{E}[T(\mathcal{G}_{n,q})] \sim \binom{n}{3}q^3$. On the other hand, we also will prove, in Proposition 3.5, the more surprising result that for some values of (p, t) the importance sampling scheme based on the edge tilt $\mathbb{Q}_n^{h_t, 0, \alpha}$ will not be asymptotically optimal. In particular, there is a subregime of the replica symmetric phase for which the edge tilt with $h = h_t$ produces a suboptimal estimator.

Regarding the triangle tilt $\mathbb{Q}_n^{h_p, \beta, \alpha}$, our main result (Prop 3.3) is a necessary and sufficient condition on the tilt parameters for the resulting importance sampling scheme to be asymptotically optimal. Moreover, optimality can be achieved by a triangle tilt for every (p, t) in the replica symmetric phase, and even for some choices of (p, t) in the replica breaking phase, as we will show in Section 4. Thus, the triangle tilt succeeds where the edge tilt fails, because the former appropriately penalizes samples with an undesired number of triangles, whereas the latter inappropriately penalizes samples with an undesired number of edges. As mentioned in the preceding paragraph, a crucial property to be expected of such an optimal triangle tilt is that samples from the tilted measure resemble the original graph $\mathcal{G}_{n,p}$ conditioned to have at least $\binom{n}{3}t^3$ triangles.

This is indeed the case for the optimal triangle tilt, thanks to Theorem 2.4.

Finally, we remark that Theorem 2.4 draws the connection between an exponential random graph and a conditioned Erdős-Rényi graph, indicating how the parameters for the two graphs must be related in order for them to resemble each other. This relationship arises from the fact that the free energy of the exponential random graph can be expressed in a variational formulation involving the LDP rate function for the conditioned Erdős-Rényi graph. For (p, t) in the replica symmetric phase, this connection has been observed by Chatterjee and Dey [6], Chatterjee and Diaconis [7], and Lubetzky and Zhao [17]. In this paper, Theorem 2.4 generalizes this connection to include parameters (p, t) in the replica breaking phase.

Organization of the paper: We start by giving precise definitions of the various constructs arising in our study in Section 2. This culminates in Theorem 2.4 that characterizes the limiting free energy of the exponential random graph model. The rest of Section 2 is devoted to drawing a connection between the exponential random graph and Erdős-Rényi random graph conditioned on an atypical number of triangles, leading to the derivation of the triangle tilts. Section 3 discusses and proves our main results on asymptotic optimality or non-optimality of the importance sampling estimators. An explicit procedure for determining the optimal triangle tilt parameters, given (p, t) , is present in Section 4 and further expanded on in Appendix A. In Section 5, we carry out numerical simulations on moderate size networks using the various proposed tilts to illustrate and compare the viability of the importance sampling schemes. Additionally, we also discuss alternative strategies for choosing the tilt measure, hybrid tilts and conditioned triangle tilts, which are variants of the edge and triangle tilts.

2 Large deviations, importance sampling and exponential random graphs

2.1 Large deviations for Erdős-Rényi random graphs

Before the proof of the main result, we start with a more detailed description of the large deviations principle for Erdős-Rényi random graphs and introduce the necessary constructs required in our proof. Chatterjee and Varadhan [8] have proved a general large deviation principle which is based on the theory of dense graph limits developed by [3] (See also Lovasz’s recent monograph, [14]). In this framework, a random graph is represented as a function $X(x, y) \in \widetilde{\mathcal{W}}$, where $\widetilde{\mathcal{W}}$ is the set of all measurable functions $f : [0, 1]^2 \rightarrow [0, 1]$ satisfying $f(x, y) = f(y, x)$. Specifically, a finite simple graph X on n vertices is represented by the function, or *graphon*,

$$X(x, y) = \sum_{\substack{i, j=1 \\ i \neq j}}^n X_{ij} \mathbf{1}_{[\frac{i-1}{n}, \frac{i}{n}) \times [\frac{j-1}{n}, \frac{j}{n})}(x, y) \in \widetilde{\mathcal{W}}. \tag{2.1}$$

Here we treat (X_{ij}) as a symmetric matrix with entries in $\{0, 1\}$ and $X_{ii} = 0$ for all i . In general, for a function $f \in \widetilde{\mathcal{W}}$, $f(x, y)$ can be interpreted as the probability of having an edge between vertices x and y . Then, we define the quotient space \mathcal{W} under the equivalence relation defined by $f \sim g$ if $f(x, y) = g(\sigma x, \sigma y)$ for some measure preserving bijection $\sigma : [0, 1] \rightarrow [0, 1]$. Intuitively, an equivalence class contains graphons that are equal after a relabelling of vertices. (See, e.g., [3, 8] for further exploration and properties of the quotient space.)

By identifying a finite graph X with its graphon representation, we can consider the probability measure $\mathbb{P}_{n,p}$ as a measure induced on \mathcal{W} supported on the finite subset of

graphons of finite graphs. For $f \in \mathcal{W}$, denote

$$\mathcal{E}(f) = \int_0^1 \int_0^1 f(x, y) dx dt \tag{2.2}$$

and

$$\mathcal{T}(f) = \int_0^1 \int_0^1 \int_0^1 f(x, y)f(y, z)f(x, z) dx dy dz. \tag{2.3}$$

We see that $E(X) = \frac{n^2}{2}\mathcal{E}(X)$ and $T(X) = \frac{n^3}{6}\mathcal{T}(X)$, so that \mathcal{E} and \mathcal{T} represent edge and triangle densities of the graph X , respectively. Then, rather than considering the event $\mathcal{W}_{n,t}$, we shall equivalently consider the upper tails of triangle densities,

$$\mathcal{W}_t := \{f \in \mathcal{W} \mid \mathcal{T}(f) \geq t^3\}.$$

The large deviation principle of Chatterjee and Varadhan [8] implies for any $p \in (0, 1)$ and $t \in [p, 1]$,

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(\mathcal{T}(\mathcal{G}_{n,p}) \geq t^3) = -\phi(p, t) \tag{2.4}$$

where $\phi(p, t)$ is the large deviation decay rate given by a variational form,

$$\phi(p, t) = \inf \{ \mathcal{I}_p(f) \mid f \in \mathcal{W}, \mathcal{T}(f) \geq t^3 \} = \inf_{f \in \mathcal{W}_t} [\mathcal{I}_p(f)]. \tag{2.5}$$

Here,

$$\mathcal{I}_p(f) := \int_0^1 \int_0^1 I_p(f(x, y)) dx dy \tag{2.6}$$

is the large deviation rate function, where $I_p : [0, 1] \rightarrow \mathbb{R}$ is defined at (1.5). A further important consequence of the large deviation principle concerns the typical behaviour of the conditioned probability measure

$$\mathbb{P}_{n,p}(A|\mathcal{W}_t) = \mathbb{P}_{n,p}(A \cap \mathcal{W}_t)\mu_n^{-1}.$$

When we refer to $\mathcal{G}_{n,p}$ conditioned on the event $\mathcal{W}_t = \{\mathcal{T}(f) \geq t^3\}$, we mean the random graph whose law is given by this conditioned probability measure.

Lemma 2.1. ([8, Theorem 3.1], Lemma C.1) *Let $\mathcal{F}^* \subset \mathcal{W}$ be the non-empty set of graphs that optimize the variational form in (2.5). Then the Erdős-Rényi graph $\mathcal{G}_{n,p}$ conditioned on $\{\mathcal{T}(f) \geq t^3\}$ is asymptotically indistinguishable from the minimal set \mathcal{F}^* .*

The term "asymptotically indistinguishable" in Lemma 2.1 roughly means that the graphon representation of the graph converges in probability, under the *cut distance* metric, to some function $f^* \in \mathcal{F}^*$ at an exponential rate as $n \rightarrow \infty$. Intuitively, this means that the typical conditioned Erdős-Rényi graph resembles some graph $f^* \in \mathcal{F}^*$ for large n . In order to give a more precise definition of asymptotic indistinguishability, we first recall the cut distance metric δ_{\square} , defined for $f, g \in \mathcal{W}$ by

$$\delta_{\square}(f, g) = \inf_{\sigma} \sup_{S, T \subset [0,1]} \left| \int_{S \times T} (f(\sigma x, \sigma y) - g(x, y)) dx dy \right|,$$

where the infimum is taken over all measure-preserving bijections $\sigma : [0, 1] \rightarrow [0, 1]$. For $\mathcal{F}_1, \mathcal{F}_2 \subset \mathcal{W}$,

$$\delta_{\square}(\mathcal{F}_1, \mathcal{F}_2) = \inf_{f_1 \in \mathcal{F}_1, f_2 \in \mathcal{F}_2} \delta_{\square}(f_1, f_2).$$

It is known by [15] that $(\mathcal{W}, \delta_{\square})$ is a compact metric space.




Figure 1: Phase diagram for $t \geq p$. The shaded region is the replica symmetric phase, and the unshaded region to the left of the curve is the replica breaking phase.

We say that a family of random graphs \mathcal{G}_n on n vertices, for $n \in \mathbb{N}$, is asymptotically indistinguishable from a subset $\mathcal{F} \subset \mathcal{W}$ if: for any $\epsilon_1 > 0$ there is $\epsilon_2 > 0$ such that

$$\limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(\delta_{\square}(\mathcal{G}_n, \mathcal{F}) > \epsilon_1) < -\epsilon_2. \quad (2.7)$$

Further, we say that \mathcal{G}_n is asymptotically indistinguishable from the *minimal* set $\mathcal{F} \subset \mathcal{W}$ if \mathcal{F} is the smallest closed subset of \mathcal{W} that \mathcal{G}_n is asymptotically indistinguishable from. Clearly, if \mathcal{G}_n is asymptotically indistinguishable from a singleton set \mathcal{F} , then \mathcal{F} is, trivially, minimal. Finally, we say two random graphs $\mathcal{G}_n^1, \mathcal{G}_n^2$ are asymptotically indistinguishable if they are each asymptotically indistinguishable from the same minimal set $\mathcal{F} \subset \mathcal{W}$. Intuitively, this means that the random behaviour, or the typical graphs, of \mathcal{G}_n^1 resembles that of \mathcal{G}_n^2 for large n . (See [7] and [8] for a wide-ranging exploration of this metric in the context of describing limits of dense random graph sequences.)

Using this terminology, we observe that an Erdős-Rényi graph $\mathcal{G}_{n,u}$ is asymptotically indistinguishable from the singleton set containing the constant function $f^* \equiv u$. A special notion about whether the conditioned Erdős-Rényi graph is again an Erdős-Rényi graph leads to the following definition.

Definition 2.2. *The replica symmetric phase is the regime of parameters (p, t) for which the large deviations rate satisfies*

$$\inf_{f \in \mathcal{W}_t} [\mathcal{I}_p(f)] = I_p(t), \quad (2.8)$$

and the infimum is uniquely attained at the constant function t .

The replica breaking phase is the regime of parameters (p, t) that are not in the replica symmetric phase. ■

The complete characterization of the replica symmetry phase was shown by [17], and is illustrated in Figure 1.

The notion of replica symmetry is a property of the rare event problem, where the Erdős-Rényi graph $\mathcal{G}_{n,p}$ conditioned on the event $\{\mathcal{T}(f) \geq t^3\}$ is asymptotically indistinguishable from an Erdős-Rényi graph with the higher edge density, $\mathcal{G}_{n,t}$, a consequence of Lemma 2.1. In contrast, the conditioned graphs in the replica breaking phase are not indistinguishable from any one Erdős-Rényi graph; instead, they may behave like

a mixture of Erdős-Rényi graphs or exhibit a clique-like structure with edge density less than t . The term “replica symmetric phase” is borrowed from [8], which in turn was inspired by the statistical physics literature. However, we remark that this term has been used differently from us by other authors to refer to other families of graphs behaving like an Erdős-Rényi graph or a mixture of Erdős-Rényi graphs.

2.2 Asymptotic behavior of exponential random graphs

To find “good” importance sampling tilted measures, we focus on the class of exponential random graphs. The exponential random graph is a random graph on n vertices defined by the Gibbs measure

$$\mathbb{Q}(X) = \mathbb{Q}_n^{h,\beta,\alpha}(X) \propto e^{n^2\mathcal{H}(X)} \tag{2.9}$$

on Ω_n , where for given $h \in \mathbb{R}$, $\beta \in \mathbb{R}_+$, $\alpha > 0$, the Hamiltonian is

$$\mathcal{H}(X) = \frac{h}{2}\mathcal{E}(X) + \frac{\beta}{6}\mathcal{T}(X)^\alpha. \tag{2.10}$$

We will use $\psi_n = \psi_n^{h,\beta,\alpha}$ to denote the log of the normalizing constant (free energy)

$$\psi_n = \psi_n^{h,\beta,\alpha} = \frac{1}{n^2} \log \sum_{X \in \Omega_n} e^{n^2\mathcal{H}(X)},$$

so that $\mathbb{Q}_n^{h,\beta,\alpha}(X) = \exp(n^2(\mathcal{H}(X) - \psi_n))$. We denote by $\mathcal{G}_n^{h,\beta,\alpha}$ the exponential random graph defined by the Gibbs measure (2.9). The case where $\alpha = 1$ is the “classical” exponential random graph model that has an enormous literature in the social sciences, see e.g. [20, 19] and the references therein and rigorously studied in a number of recent papers, see e.g. [1, 7, 18, 17, 22, 23]. In this case, the Hamiltonian can be rewritten as $n^2\mathcal{H}(X) = hE(X) + \frac{\beta}{n}T(X)$. We will drop the superscripts in $\psi_n^{h,\beta}$, $\mathbb{Q}_n^{h,\beta}$ when $\alpha = 1$. The generalization to the exponential random graph with the parameter α was first proposed in [17].

Observe that the Erdős-Rényi random graph is a special case of the exponential random graph: if $\beta = 0$ and $h = h_p$ with h_p defined by (1.2), then $\mathbb{Q}_n^{h_p,0,\alpha} = \mathbb{P}_{n,p}$ for any $\alpha > 0$ and the edges are independent with probability p . On the other hand, choosing $\beta > 0$ introduces a non-trivial dependence between the edges. By adjusting the parameters (h, β, α) , the Gibbs measure $\mathbb{Q}_n^{h,\beta,\alpha}$ can be adjusted to favor edges and triangles to varying degree.

The asymptotic behavior of the exponential random graph measures $\mathbb{Q}_n^{h,\beta,\alpha}$ and the free energy $\psi_n^{h,\beta,\alpha}$ is partially characterized by the following result of Chatterjee and Diaconis [7] and Lubetzky and Zhao [17]. In what follows, we will make use of the functions

$$I(u) = \frac{1}{2}u \log u + \frac{1}{2}(1 - u) \log(1 - u) \tag{2.11}$$

on $u \in [0, 1]$ and, for $f \in \mathcal{W}$,

$$\mathcal{I}(f) := \int_0^1 \int_0^1 I(f(x, y)) dx dy. \tag{2.12}$$

Theorem 2.3 (See [7] [17]). *For the exponential random graph $\mathcal{G}_n^{h,\beta,\alpha}$ with parameters $(h, \beta, \alpha) \in \mathbb{R} \times [0, +\infty) \times [2/3, 1]$, the free energy satisfies*

$$\lim_{n \rightarrow \infty} \psi_n^{h,\beta,\alpha} = \sup_{0 \leq u \leq 1} \left[\frac{\beta}{6}u^{3\alpha} - I(u) + \frac{h}{2}u \right]. \tag{2.13}$$

If the supremum in (2.13) is attained at a unique point $v^ \in [0, 1]$, then the exponential random graph $\mathcal{G}_n^{h,\beta,\alpha}$ is asymptotically indistinguishable from the Erdős-Rényi graph \mathcal{G}_{n,v^*} .*

The case $\alpha = 1$ in Theorem 2.3 was proved by Chatterjee and Diaconis – Theorems 4.1, 4.2 of [7]; the cases $\alpha \in [2/3, 1]$ is due to Lubetzky and Zhao– Theorems 1.3, 4.3 of [17].

Our main result in this section, stated next, is the generalization of the variational formulation for the free energy of the Gibbs measure of any exponential random graph. Our result emphasizes the connection between the exponential random graph and the conditioned Erdős-Rényi graph. Before stating the result we will need some extra notation. Extend the Hamiltonian defined in (2.10) to the space of graphons in the natural way

$$\mathcal{H}(f) := \frac{h}{2}\mathcal{E}(f) + \frac{\beta}{6}\mathcal{T}(f)^\alpha \tag{2.14}$$

where recall the definitions for the density of edges and triangles for graphons defined respectively in (2.2) and (2.3). For fixed $q \in (0, 1)$ recall the functions $\mathcal{I}_q(f)$ from (2.6) and the function $\mathcal{I}(f)$ from (2.12). In particular, observe that

$$\mathcal{I}_q(f) = \mathcal{I}(f) - \frac{h_q}{2}\mathcal{E}(f) - \frac{1}{2}\log(1 - q). \tag{2.15}$$

with $h_q = \log \frac{q}{1-q}$.

Theorem 2.4. *For the exponential random graph $\mathcal{G}_n^{h,\beta,\alpha}$ with parameters $(h, \beta, \alpha) \in \mathbb{R} \times [0, +\infty) \times [0, 1]$, the free energy satisfies*

$$\lim_{n \rightarrow \infty} \psi_n^{h,\beta,\alpha} = \sup_{0 \leq u \leq 1} \left[\frac{\beta}{6}u^{3\alpha} - \phi_q(u) - \frac{1}{2}\log(1 - q) \right] \tag{2.16}$$

where $q \in (0, 1)$ is such that $h = h_q = \log \frac{q}{1-q}$, and

$$\phi_q(u) = \inf_{f \in \partial\mathcal{W}_u} [\mathcal{I}_q(f)] \tag{2.17}$$

and $\partial\mathcal{W}_u := \{f \in \mathcal{W} \mid \mathcal{T}(f) = u^3\}$.

If the supremum in (2.16) is attained at a unique point $v^* \geq q$, then the exponential random graph $\mathcal{G}_n^{h_q,\beta,\alpha}$ is asymptotically indistinguishable from the conditioned Erdős-Rényi graph, $\mathcal{G}_{n,q}$ conditioned on the event $\{\mathcal{T}(f) \geq (v^*)^3\}$.

Remark 2.5. [(i)]

1. If, in addition, (q, v^*) in Theorem 2.4 belongs to the replica symmetric phase, then $\mathcal{G}_n^{h_q,\beta,\alpha}$ is asymptotically indistinguishable from the Erdős-Rényi graph \mathcal{G}_{n,v^*} . This follows from the remarks following Definition 2.2, that in the replica symmetric phase, $\mathcal{G}_{n,q}$ conditioned on $\{\mathcal{T}(f) \geq (v^*)^3\}$ is asymptotically indistinguishable from the Erdős-Rényi graph \mathcal{G}_{n,v^*} . In this case, (2.16) reduces to (2.13).
2. Non-uniqueness of v^* is possible. As will be apparent from the proof, if the supremum in (2.16) is attained on the set $U^* \subset [0, 1]$, then the exponential random graph $\mathcal{G}_n^{h_q,\beta,\alpha}$ is asymptotically indistinguishable from the minimal set $\mathcal{F}^* = \bigcup_{u \in U^*} \mathcal{F}_u^*$, where \mathcal{F}_u^* is the set of minimizers of (2.17). In particular, if U^* contains more than one element, then $\mathcal{G}_n^{h_q,\beta,\alpha}$ is asymptotically indistinguishable from a mixture of different conditioned Erdős-Rényi graphs.

Proof. Theorem 3.1 in [7] implies that

$$\lim_{n \rightarrow \infty} \psi_n^{h_q,\beta,\alpha} = \sup_{f \in \mathcal{W}} [\mathcal{H}(f) - \mathcal{I}(f)]. \tag{2.18}$$

To show (2.16), suppose $f \in \partial\mathcal{W}_u$, for $u \in (0, 1)$. Recalling (2.15), we have

$$\begin{aligned} \mathcal{H}(f) - \mathcal{I}(f) &= \frac{h_q}{2} \mathcal{E}(f) + \frac{\beta}{6} u^{3\alpha} - \mathcal{I}(f) \\ &= \frac{\beta}{6} u^{3\alpha} - \mathcal{I}_q(f) - \frac{1}{2} \log(1 - q) \\ &\leq \frac{\beta}{6} u^{3\alpha} - \inf_{f \in \partial\mathcal{W}_u} [\mathcal{I}_q(f)] - \frac{1}{2} \log(1 - q) \end{aligned} \tag{2.19}$$

This implies that

$$\sup_{f \in \partial\mathcal{W}_u} [\mathcal{H}(f) - \mathcal{I}(f)] \leq \frac{\beta}{6} u^{3\alpha} - \inf_{f \in \partial\mathcal{W}_u} [\mathcal{I}_q(f)] - \frac{1}{2} \log(1 - q),$$

and

$$\begin{aligned} \sup_{f \in \mathcal{W}} [\mathcal{H}(f) - \mathcal{I}(f)] &= \sup_{0 \leq u \leq 1} \sup_{f \in \partial\mathcal{W}_u} [\mathcal{H}(f) - \mathcal{I}(f)] \\ &\leq \sup_{0 \leq u \leq 1} \left[\frac{\beta}{6} u^{3\alpha} - \inf_{f \in \partial\mathcal{W}_u} [\mathcal{I}_q(f)] - \frac{1}{2} \log(1 - q) \right] \end{aligned}$$

Now we show the reverse inequality. Fix $\epsilon > 0$. For each $u \in (0, 1)$, let $f_{u,\epsilon} \in \partial\mathcal{W}_u$ be such that

$$\mathcal{I}_q(f_{u,\epsilon}) \leq \inf_{f \in \partial\mathcal{W}_u} [\mathcal{I}_q(f)] + \epsilon.$$

Therefore, for each $u \in (0, 1)$ we have

$$\begin{aligned} \mathcal{H}(f_{u,\epsilon}) - \mathcal{I}(f_{u,\epsilon}) &= \frac{\beta}{6} u^{3\alpha} - [\mathcal{I}_q(f_{u,\epsilon})] - \frac{1}{2} \log(1 - q) \\ &\geq \frac{\beta}{6} u^{3\alpha} - \inf_{f \in \partial\mathcal{W}_u} [\mathcal{I}_q(f)] - \frac{1}{2} \log(1 - q) - \epsilon. \end{aligned} \tag{2.20}$$

Hence

$$\sup_{f \in \mathcal{W}} [\mathcal{H}(f) - \mathcal{I}(f)] \geq \sup_{0 < u < 1} \left[\frac{\beta}{6} u^{3\alpha} - \inf_{f \in \partial\mathcal{W}_u} [\mathcal{I}_q(f)] - \frac{1}{2} \log(1 - q) \right] - \epsilon. \tag{2.21}$$

Since $\epsilon > 0$ is arbitrary, (2.16) follows.

To show the next statement, suppose the supremum in (2.16) is attained at a unique point $v^* \geq q$. Let $\mathcal{F}_{v^*}^* \subset \partial\mathcal{W}_{v^*}$ denote the set of functions that attains the infimum in (2.17). We observe from the preceding proof that, in fact, $\mathcal{F}_{v^*}^*$ is the set that attains the infimum in (2.18), so that by [7, Theorem 3.2] and Lemma C.1, the graph $\mathcal{G}_n^{h_q, \beta, \alpha}$ is asymptotically indistinguishable from the minimal set $\mathcal{F}_{v^*}^*$. On the other hand, since $\mathcal{F}_{v^*}^*$ is also the set that attains the infimum in the LDP rate in (2.5) (due to [8, Theorem 4.2(iii)]), the conditioned Erdős-Rényi graph, $\mathcal{G}_{n,q}$ conditioned on $\{\mathcal{T}(f) \geq (v^*)^3\}$, is also asymptotically indistinguishable from the set $\mathcal{F}_{v^*}^*$. Thus, $\mathcal{G}_n^{h_q, \beta, \alpha}$ is asymptotically indistinguishable from the conditioned Erdős-Rényi graph, $\mathcal{G}_{n,q}$ conditioned on the event $\{\mathcal{T}(f) \geq (v^*)^3\}$. ■

The mean behaviour of the triangle density of an exponential random graph $\mathcal{G}_n^{h_q, \beta, \alpha}$ can be deduced from the variational formulation in (2.16), and in special instances, so can the mean behaviour of the edge density. This is shown in the next proposition, which follows from [7, Theorem 4.2] and the Lipschitz continuity of the mappings $f \mapsto \mathcal{T}(f)$ and $f \mapsto \mathcal{E}(f)$ under the cut distance metric δ_\square [3, Theorem 3.7]. The proof is left to the appendix.

Proposition 2.6. *Let $(h_q, \beta, \alpha) \in \mathbb{R} \times [0, +\infty) \times [0, 1]$. If the supremum in (2.16) is attained at a unique point $v^* \in [0, 1]$, then*

$$\lim_{n \rightarrow \infty} \mathbb{E} |\mathcal{T}(\mathcal{G}_n^{h_q, \beta, \alpha}) - (v^*)^3| = 0. \tag{2.22}$$

Further, if (q, v^*) belongs to the replica symmetric phase, then

$$\lim_{n \rightarrow \infty} \mathbb{E} |\mathcal{E}(\mathcal{G}_n^{h_q, \beta, \alpha}) - v^*| = 0. \tag{2.23}$$

3 Asymptotic Optimality

Recall that the edge tilt corresponds to the Gibbs measure (1.13) with $\beta = 0$ and $h > h_p = \log \frac{p}{1-p}$. Thus, an edge tilt $\mathbb{Q}_n^{h,0}$ satisfies

$$\frac{d\mathbb{P}_{n,p}}{d\mathbb{Q}_n^{h,0}}(X) = \exp \left[-n^2 \left(\frac{h - h_p}{2} \mathcal{E}(X) + \psi_n^{h_p,0} - \psi_n^{h,0} \right) \right]. \tag{3.1}$$

The triangle tilt corresponds to the Gibbs measure (1.13) with $h = h_p$ and $\beta > 0, \alpha > 0$. So, the triangle tilt $\mathbb{Q}_n^{h_p, \beta, \alpha}$ satisfies

$$\frac{d\mathbb{P}_{n,p}}{d\mathbb{Q}_n^{h_p, \beta, \alpha}}(X) = \exp \left[-n^2 \left(\frac{\beta}{6} \mathcal{T}(X)^\alpha + \psi_n^{h_p,0} - \psi_n^{h_p, \beta, \alpha} \right) \right]$$

Here recall that $\mathcal{T}(X) = \frac{6}{n^3} T(X)$ is the density of triangles in X and $\mathcal{E}(X) = \frac{2}{n^2} E(X)$ is the density of edges.

For any admissible parameters (h, β, α) , the importance sampling estimator based on the tilted measure $\mathbb{Q}_n^{h, \beta, \alpha}$ is

$$\begin{aligned} \tilde{M}_K &= \frac{1}{K} \sum_{k=1}^K \mathbf{1}_{\mathcal{W}_t}(\tilde{X}_k) \frac{d\mathbb{P}_{n,p}}{d\mathbb{Q}_n^{h, \beta, \alpha}}(\tilde{X}_k) \\ &= \frac{1}{K} \sum_{k=1}^K \mathbf{1}_{\mathcal{W}_t}(\tilde{X}_k) \exp \left\{ n^2 \left(\frac{h_p - h}{2} \mathcal{E}(\tilde{X}_k) - \frac{\beta}{6} \mathcal{T}(\tilde{X}_k)^\alpha + \psi_n^{h_p,0} - \psi_n^{h, \beta, \alpha} \right) \right\} \end{aligned} \tag{3.2}$$

where \tilde{X}_k are i.i.d. samples drawn from $\mathbb{Q}_n^{h, \beta, \alpha}$. Denote

$$\hat{q}_n = \hat{q}_n(\tilde{X}) = \mathbf{1}_{\mathcal{W}_t}(\tilde{X}) \frac{d\mathbb{P}_{n,p}}{d\mathbb{Q}_n^{h, \beta, \alpha}}(\tilde{X}).$$

For any (h, β, α) , $\mathbb{E}[\hat{q}_n] = \mu_n$ and so \tilde{M}_K is an unbiased estimator for μ_n .

Our first result is a necessary condition for asymptotic optimality of the importance sampling scheme:

Proposition 3.1. *Given $p < t$, let $(h, \beta, \alpha) \in \mathbb{R} \times [0, +\infty) \times [0, 1]$ with $h = h_q = \log \frac{q}{1-q}$. Suppose that the supremum in (2.16) is not attained at t :*

$$\sup_{0 \leq u \leq 1} \left[\frac{\beta}{6} u^{3\alpha} - \phi_q(u) - \frac{1}{2} \log(1 - q) \right] \neq \frac{\beta}{6} t^{3\alpha} - \phi_q(t) - \frac{1}{2} \log(1 - q). \tag{3.3}$$

Then the importance sampling scheme based on the Gibbs measure tilt $\mathbb{Q}_n^{h, \beta, \alpha}$ is not asymptotically optimal.

Corollary 3.2. *Given $p < t$, let $(h, \beta, \alpha) \in \mathbb{R} \times [0, +\infty) \times [0, 1]$ with $h = h_q = \log \frac{q}{1-q}$. Suppose that family of random graphs $\mathcal{G}_n^{h, \beta, \alpha}$ is not indistinguishable from $\mathcal{G}_{n,p}$ conditioned on the event $\{\mathcal{T}(X) \geq t^3\}$. Then the importance sampling scheme based on the Gibbs measure tilt $\mathbb{Q}_n^{h, \beta, \alpha}$ is not asymptotically optimal.*

Our next result shows that for triangle tilts (i.e. $h = h_p$), the necessary condition described in Proposition 3.1 is also a sufficient condition for asymptotic optimality:

Proposition 3.3. *Given $p < t$, let $h = h_p$ and $(\beta, \alpha) \in [0, +\infty) \times [0, 1]$. Suppose that the supremum in (2.16) is attained at t :*

$$\sup_{0 \leq u \leq 1} \left[\frac{\beta}{6} u^{3\alpha} - \phi_p(u) - \frac{1}{2} \log(1 - p) \right] = \frac{\beta}{6} t^{3\alpha} - \phi_p(t) - \frac{1}{2} \log(1 - p). \quad (3.4)$$

Then the importance sampling scheme based on the triangle tilt $\mathbb{Q}_n^{h_p, \beta, \alpha}$ is asymptotically optimal.

In Section 4, we give a more explicit way to determine the tilt parameters that satisfy the condition (3.4).

Next, we turn to the edge tilts (i.e. $\beta = 0$). Since $\phi_q(u)$ is minimized at $u = q$ and (3.3) holds if $\beta = 0$ and $q \neq t$, we have, as a corollary of Prop 3.1, the following necessary condition for an edge tilt to produce an optimal scheme.

Proposition 3.4. *Given $p < t$, let $\beta = 0$ and $h = h_q$ for some $q \neq t$. The importance sampling scheme based on the edge tilt $\mathbb{Q}_n^{h_q, 0, \alpha}$ is not asymptotically optimal.*

Observe that for the edge tilt with $h = h_t$ and $\beta = 0$, the supremum in (2.16) is always attained at t , since

$$\inf_{0 \leq u \leq 1} [\phi_t(u)] = \inf_{f \in \mathcal{W}} [\mathcal{I}_t(f)] = \phi_t(t) = 0. \quad (3.5)$$

Thus, the edge tilt with $h = h_t$ always satisfies the necessary condition for asymptotic optimality of the importance sampling scheme. Furthermore, if (p, t) is in the replica symmetric phase, the tilted measure $\mathbb{Q}_n^{h_t, 0, \alpha}$ is indistinguishable from the conditioned Erdős-Rényi graph. Nevertheless, the sampling scheme based on the edge tilt with $h = h_t$ may still be suboptimal, even in the replica symmetric phase, as the next result shows.

Proposition 3.5. *Let $0 < p < \frac{e^{-1/2}}{1+e^{-1/2}}$ and $t \in (p, 1)$. If t is sufficiently close to 1 and (p, t) belong to the replica symmetric phase, then the importance sampling scheme based on the edge tilt $\mathbb{Q}_n^{h_t, 0}$ is not asymptotically optimal.*

Remark A.5 and Figure 3 indicate that there do exist parameters (p, t) belonging to the replica symmetric phase for which the hypothesis of Prop 3.5 is satisfied.

3.1 Proofs of results.

We first prove the asymptotic optimality of the triangle tilts, Prop 3.3.

Proof of Proposition 3.3. Due to (1.10), it suffices to show that

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{\mathbb{Q}}[\hat{q}_n^2] \leq -2 \inf_{f \in \mathcal{W}_t} \mathcal{I}_p(f). \quad (3.6)$$

Let $q \in (0, 1)$ be such that $h = h_q = \log \frac{q}{1-q}$. Recall that

$$\mathbb{Q}_n^{h, \beta, \alpha}(X) = \exp \left[n^2 \left(\frac{h}{2} \mathcal{E}(X) + \frac{\beta}{6} \mathcal{T}^\alpha(X) - \psi_n^{h, \beta, \alpha} \right) \right].$$

Therefore, by definition of \hat{q}_n , we have

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}_n}[\hat{q}_n^2] &= \mathbb{E}_{\mathbb{P}_{n,p}} \left[\mathbf{1}_{\mathcal{W}_t} \frac{d\mathbb{P}_{n,p}}{d\mathbb{Q}_n^{h_q, \beta, \alpha}} \right] \\ &= \mathbb{E}_{\mathbb{P}_{n,p}} \left[\mathbf{1}_{\mathcal{W}_t} \exp \left\{ n^2 \left(\frac{h_p - h_q}{2} \mathcal{E}(X) - \frac{\beta}{6} \mathcal{T}(X)^\alpha + \psi_n^{h_q, \beta, \alpha} - \psi_n^{h_p, 0} \right) \right\} \right]. \end{aligned}$$

The mappings $\mathcal{E}, \mathcal{T} : \mathcal{W} \mapsto \mathbb{R}$ are bounded and continuous [3, Theorem 3.8]. Applying the Laplace principle for the family of measures $\mathbb{P}_{n,p}$, for which $\mathcal{I}_p(f)$ is the rate function [7, Theorem 3.1], we obtain

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{\mathbb{Q}_n} [\hat{q}_n^2] &= \lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{\mathbb{P}_{n,p}} \left[\mathbf{1}_{\mathcal{W}_t} \exp \left\{ n^2 \left(\frac{h_p - h_q}{2} \mathcal{E}(X) - \frac{\beta}{6} \mathcal{T}(X)^\alpha \right) \right\} \right] \\ &\quad + \lim_{n \rightarrow \infty} (\psi_n^{h_q, \beta, \alpha} - \psi_n^{h_p, 0}) \\ &= - \inf_{f \in \mathcal{W}_t} \left[\mathcal{I}_p(f) + \frac{h_q - h_p}{2} \mathcal{E}(f) + \frac{\beta}{6} \mathcal{T}(f)^\alpha \right] + \lim_{n \rightarrow \infty} (\psi_n^{h_q, \beta, \alpha} - \psi_n^{h_p, 0}) \end{aligned} \tag{3.7}$$

By (2.16),

$$\lim_{n \rightarrow \infty} \psi_n^{h_q, \beta, \alpha} = V(u^*)$$

where $u^* = \operatorname{argsup}_{0 \leq u \leq 1} [V(u)]$ and

$$V(u) := \frac{\beta}{6} u^{3\alpha} - \inf_{f \in \partial \mathcal{W}_u} [\mathcal{I}_q(f)] - \frac{1}{2} \log(1 - q).$$

Also, $\lim_{n \rightarrow \infty} \psi_n^{h_p, 0} = -\frac{1}{2} \log(1 - p)$. Hence,

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{\mathbb{Q}} [\hat{q}_n^2] &\tag{3.8} \\ &= - \inf_{f \in \mathcal{W}_t} \left[\mathcal{I}_p(f) + \frac{h_q - h_p}{2} \mathcal{E}(f) + \frac{\beta}{6} \mathcal{T}(f)^\alpha \right] + \frac{\beta}{6} (u^*)^{3\alpha} - \inf_{f \in \partial \mathcal{W}_{u^*}} [\mathcal{I}_q(f)] - \frac{1}{2} \log \frac{1 - q}{1 - p} \\ &\leq - \inf_{f \in \mathcal{W}_t} \left[\mathcal{I}_p(f) + \frac{h_q - h_p}{2} \mathcal{E}(f) \right] + \frac{\beta}{6} ((u^*)^{3\alpha} - t^{3\alpha}) - \inf_{f \in \partial \mathcal{W}_{u^*}} [\mathcal{I}_q(f)] - \frac{1}{2} \log \frac{1 - q}{1 - p} \end{aligned}$$

The last inequality follows from the fact that $\mathcal{T}(f) \geq t^3$ for all $f \in \mathcal{W}_t$. Since,

$$\mathcal{I}_q(f) + \frac{1}{2} \log \frac{1 - q}{1 - p} = \mathcal{I}_p(f) + \frac{h_p - h_q}{2} \mathcal{E}(f),$$

we conclude that

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{\mathbb{Q}} [\hat{q}_n^2] &\leq - \inf_{f \in \mathcal{W}_t} \left[\mathcal{I}_p(f) + \frac{h_q - h_p}{2} \mathcal{E}(f) \right] - \inf_{f \in \partial \mathcal{W}_{u^*}} \left[\mathcal{I}_p(f) - \frac{h_q - h_p}{2} \mathcal{E}(f) \right] \\ &\quad + \frac{\beta}{6} ((u^*)^{3\alpha} - t^{3\alpha}) \end{aligned} \tag{3.9}$$

The estimate (3.9) holds for any $(h_q, \beta, \alpha) \in \mathbb{R} \times [0, +\infty) \times [0, 1]$. However, under the hypotheses of Proposition 3.3, we have $u^* = t$ and $q = p$. Therefore,

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{\mathbb{Q}} [\hat{q}_n^2] \leq - \inf_{f \in \mathcal{W}_t} [\mathcal{I}_p(f)] - \inf_{f \in \partial \mathcal{W}_t} [\mathcal{I}_p(f)] = -2 \inf_{f \in \mathcal{W}_t} [\mathcal{I}_p(f)]$$

Combined with the upper bound for the asymptotic second moment, we conclude that the triangle tilt $\mathbb{Q}_n^{h_p, \beta, \alpha}$ yields an asymptotically optimal importance sampling estimator if (3.4) holds. ■

We now prove the necessary condition for optimality, Prop 3.1.

Proof of Proposition 3.1. We recall from (2.18) that $\lim_{n \rightarrow \infty} \psi_n^{h_q, \beta, \alpha} = \sup_{f \in \mathcal{W}} [\mathcal{H}(f) - \mathcal{I}(f)]$. Due to Theorem 2.4, there exists $f^* \in \mathcal{W}$ such that f^* minimizes the LDP rate function $\inf_{f \in \mathcal{W}_t} [\mathcal{I}_p(f)]$, and f^* does not maximize $\sup_{f \in \mathcal{W}} [\mathcal{H}(f) - \mathcal{I}(f)]$. From (3.7),

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{\mathbb{Q}_n} [\hat{q}_n^2] \\ &= - \inf_{f \in \mathcal{W}_t} \left[\mathcal{I}_p(f) + \frac{h - h_p}{2} \mathcal{E}(f) + \frac{\beta}{6} \mathcal{T}(f)^\alpha \right] + \lim_{n \rightarrow \infty} \psi_n^{h, \beta, \alpha} + \frac{1}{2} \log(1 - p) \\ &= - \inf_{f \in \mathcal{W}_t} \left[\mathcal{I}_p(f) + \frac{h - h_p}{2} \mathcal{E}(f) + \frac{\beta}{6} \mathcal{T}(f)^\alpha \right] + \sup_{f \in \mathcal{W}} [\mathcal{H}(f) - \mathcal{I}(f)] + \frac{1}{2} \log(1 - p) \\ &> - \left[\mathcal{I}_p(f^*) + \frac{h - h_p}{2} \mathcal{E}(f^*) + \frac{\beta}{6} \mathcal{T}(f^*)^\alpha \right] + \frac{h}{2} \mathcal{E}(f^*) + \frac{\beta}{6} \mathcal{T}(f^*)^\alpha - \mathcal{I}(f^*) + \frac{1}{2} \log(1 - p) \\ &= -2\mathcal{I}_p(f^*) = -2 \inf_{f \in \mathcal{W}_t} [\mathcal{I}_p(f)] \end{aligned}$$

Hence the importance sampling estimator is not asymptotically optimal. ■

Proof of Proposition 3.5. For the edge tilt with $h = h_t, \beta = 0$, we have from (3.8),

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{\mathbb{Q}} [\hat{q}_n^2] = - \inf_{f \in \mathcal{W}_t} \left[\mathcal{I}_p(f) + \left(\frac{h_t - h_p}{2} \right) \mathcal{E}(f) \right] - I_p(t) + \left(\frac{h_t - h_p}{2} \right) t. \quad (3.10)$$

Because (p, t) is in the replica symmetric phase, the term $\mathcal{I}_p(f)$ is minimized over \mathcal{W}_t by the constant function,

$$f_t(x, y) \equiv t = \arg \inf_{f \in \mathcal{W}_t} [\mathcal{I}_p(f)].$$

On the other hand, the term $\mathcal{E}(f)$ is minimized over \mathcal{W}_t by the clique function

$$g_t(x, y) = \mathbf{1}_{[0, t]^2}(x, y) = \arg \inf_{f \in \mathcal{W}_t} [\mathcal{E}(f)]. \quad (3.11)$$

This g_t represents a graph with a large clique, in which there is a complete subgraph on a fraction t of the vertices. Let $\mathcal{V}(f) = \mathcal{I}_p(f) + \frac{h_t - h_p}{2} \mathcal{E}(f)$ be the function to be minimized in (3.10). We have

$$\mathcal{V}(f_t) = I_p(t) + \left(\frac{h_t - h_p}{2} \right) t,$$

and

$$\mathcal{V}(g_t) = t^2 I_p(1) + (1 - t^2) I_p(0) + \left(\frac{h_t - h_p}{2} \right) t^2.$$

Thus, if we can show that

$$\mathcal{V}(g_t) < \mathcal{V}(f_t),$$

it will follow that from (3.10) that

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{\mathbb{Q}} [\hat{q}_n^2] &\geq -\mathcal{V}(g_t) - I_p(t) + \left(\frac{h_t - h_p}{2} \right) t \\ &> -2I_p(t) = -2 \inf_{f \in \mathcal{W}_t} \mathcal{I}_p(f). \end{aligned} \quad (3.12)$$

We claim that for $p < \frac{e^{-1/2}}{1 + e^{-1/2}}$ and t sufficiently close to 1, we have $\mathcal{V}(g_t) < \mathcal{V}(f_t)$. Indeed, let

$$\begin{aligned} G(t) &:= \mathcal{V}(g_t) - \mathcal{V}(f_t) \\ &= t^2 I_p(1) + (1 - t^2) I_p(0) - I_p(t) + \left(\frac{h_t - h_p}{2} \right) (t^2 - t). \end{aligned} \quad (3.13)$$

Observe that $G(1) = 0$ and

$$G'(1) = 2I_p(1) - 2I_p(0) - 1/2 = -\log\left(\frac{p}{1-p}\right) - 1/2.$$

So, $G'(1) > 0$ if $h_p < -1/2$, i.e., if $p < \frac{e^{-1/2}}{1+e^{-1/2}}$. So, for t sufficiently close to 1, we have $\mathcal{V}(g_t) < \mathcal{V}(f_t)$, and we conclude that (3.12) holds with strict inequality. Hence, the importance sampling scheme associated with the edge tilt $Q_n^{h_t, 0}$ cannot be asymptotically optimal. ■

4 Characterizing regimes for the triangle tilt

Proposition 3.3 describes the necessary and sufficient condition (3.4) on the parameters (β, α) of a triangle tilt, that will produce an optimal importance sampling scheme. Given (p, t) , do these optimal tilt parameters (β, α) exist and how can they be found? In this section, we describe in a pseudo-explicit procedure for determining the optimal tilt parameters given (p, t) .

An explicit determination of the optimal tilt parameters can be made when (p, t) belongs to the replica symmetry phase.

Proposition 4.1. *If (p, t) belongs to the replica symmetry phase, then there exists some $\alpha \in [2/3, 1]$ for which the triangle tilt with parameters (h_p, β, α) produces an optimal scheme, where β satisfies*

$$\beta = \frac{h_t - h_p}{\alpha t^{3\alpha-1}}. \tag{4.1}$$

It will turn out that if there exists some $\alpha \in [0, 1]$ and some β for which the triangle tilt produces an optimal scheme, then for any $\alpha' \in [0, \alpha]$, and an appropriate β' depending on α' , the triangle tilt with parameters (h_p, β', α') also produces an optimal scheme (see Lemma A.4). Thus, in Prop 4.1, we can always take $\alpha = 2/3$.

It is also of interest to determine the tilt parameters when (p, t) belong to the replica breaking phase. Our next result, Prop 4.3, states a more general characterization of the optimal tilt parameters that applies to both the replica symmetry and breaking phases. To state the result, we introduce the minorant condition.

We shall say that (p, t) satisfies the *minorant condition with parameter α* if the point $(t^{3\alpha}, \phi_p(t))$ lies on the convex minorant of the function $x \mapsto \phi_p(x^{1/3\alpha})$. In this case, subdifferential(s) of the convex minorant of $x \mapsto \phi_p(x^{1/3\alpha})$ at $x = t^{3\alpha}$ always exist and are positive. Recall that the subdifferentials of a convex function $f(x)$ at a point x are the slopes of any line lying below $f(x)$ that is tangent to f at x . The set of subdifferentials of a convex function is non-empty; if the function is differentiable at x , then the set of subdifferentials contains exactly one point, the derivative $f'(x)$.

The minorant condition is not an unattainable one, as shown in the next lemma.

Lemma 4.2. *The parameters (p, t) that satisfy the minorant condition with some α includes the replica symmetry phase as well as a non-empty subset of the replica breaking phase.*

Proposition 4.3. *Suppose (p, t) satisfies the minorant condition for some $\alpha \in [0, 1]$. Then the triangle tilt with the parameters (h_p, β, α) produces an optimal scheme, where β is such that $\frac{\beta}{6}$ is a subdifferential of the convex minorant of $x \mapsto \phi_p(x^{1/3\alpha})$ at $x = t^{3\alpha}$.*

Moreover, if $\phi_p(u)$ is differentiable at t , then

$$\beta = \frac{2\phi'_p(t)}{\alpha t^{3\alpha-1}}. \tag{4.2}$$

Combining Lemma 4.2 and Prop 4.3, there exists (p, t) belonging to the replica breaking phase for which a triangle tilt that produces an optimal scheme exists. In particular, in the replica symmetry phase, since $\phi_p(t) = I_p(t)$ is differentiable at t , Prop 4.3 reduces to Prop 4.1. Thus, Prop 4.1 gives an explicit construction of the tilt parameters when (p, t) belong to the replica symmetry phase. In the replica breaking phase, we may need to resort to numerical strategies to find the tilt parameters. Nonetheless, we emphasize that it is possible in principle to construct an optimal importance sampling estimator in the replica breaking phase even if the limiting behaviour of the conditioned graph is not known exactly.

Proofs of results.

Notice that if t attains the supremum in (3.4), we may rewrite the condition as

$$t = \arg \sup_{0 \leq u \leq 1} \left[\frac{\beta}{6} u^{3\alpha} - \phi_p(u) \right].$$

Together with Prop 3.3, the next lemma immediately implies Prop 4.3.

Lemma 4.4. *Suppose (p, t) satisfies the minorant condition for some $\alpha > 0$. Let β be such that $\frac{\beta}{6}$ is a subdifferential of the convex minorant of $x \mapsto \phi_p(x^{1/3\alpha})$ at $x = t^{3\alpha}$. Then $\sup_{0 \leq u \leq 1} [\frac{\beta}{6} u^{3\alpha} - \phi_p(u)]$ is maximized at t .*

Moreover, if $\phi_p(u)$ is differentiable at t , then β is defined in (4.2).

Proof. The proof follows a similar technique to [17]. Using the rescaling $u \mapsto x^{1/3\alpha}$, the variational form $\sup_{0 \leq u \leq 1} [\frac{\beta}{6} u^{3\alpha} - \phi_p(u)]$ can be rewritten as

$$\sup_{0 \leq x \leq 1} \left[\frac{\beta}{6} x - \phi_p(x^{1/3\alpha}) \right].$$

Let $\hat{\phi}_p(x)$ denote the convex minorant of $x \mapsto \phi_p(x^{1/3\alpha})$. The assumption that $\frac{\beta}{6}$ is a subdifferential of $\hat{\phi}_p(x)$ at $x = t^{3\alpha}$ implies that the maximum of $\sup_x [\frac{\beta}{6} x - \hat{\phi}_p(x)]$ is attained at $t^{3\alpha}$. By the hypothesis of the lemma, we have assumed that (p, t) satisfies the minorant condition for α , so that the point $(t^{3\alpha}, \phi_p(t))$ lies on $\hat{\phi}_p(x)$. Thus we have that $\hat{\phi}_p(t^{3\alpha}) = \phi_p(t)$ and so the maximum of $\sup_x [\frac{\beta}{6} x - \phi_p(x^{1/3\alpha})]$ is also attained at $t^{3\alpha}$. It follows that the maximum of $\sup_u [\frac{\beta}{6} u^{3\alpha} - \phi_p(u)]$ is attained at t . (However, this maximum may not be unique. If the subtangent line defined by the subdifferential $\frac{\beta}{6}$ touches $\hat{\phi}_p$ at another point $r^{3\alpha}$, then r also a maximum.)

To prove the last part of the lemma, if $\phi_p(u)$ is differentiable at t , then the subdifferential is simply the derivative. Then we have

$$0 = \frac{\partial}{\partial x} \Big|_{x=t^{3\alpha}} \left[\frac{\beta}{6} x - \phi_p(x^{1/3\alpha}) \right] = \frac{\beta}{6} - \phi'_p(t) \frac{t^{1-3\alpha}}{3\alpha}$$

implies that $\beta = \frac{2\phi'_p(t)}{\alpha t^{3\alpha-1}}$. ■

Proof. (Proof of Lemma 4.2.) Recalling Definition 2.2 of the replica symmetric phase, the it follows from [17, Theorems 1.1 and 1.3] that any (p, t) that belongs to the replica symmetric phase satisfies the minorant condition for some $\alpha \in [2/3, 1]$.

We now show that there exists (p, t) belonging to the replica breaking phase that satisfies the minorant condition for some α . Notice from Appendix A and Figure 3 that there exists some a critical value p_{crit} such that when $p \geq p_{crit}$, (p, t) is replica symmetric for all $t \in [p, 1]$; whereas when $p \leq p_{crit}$, there exists an interval $(\underline{r}_p, \bar{r}_p) \subset (p, 1)$ where (p, t) is replica breaking if $t \in (\underline{r}_p, \bar{r}_p)$, and (p, t) is replica symmetric for all other values of t .

To see this, consider $\alpha = 1/3$ and convex minorant of $x \mapsto \phi_p(x^{1/3\alpha}) = \phi_p(x)$. For each $p < p_{crit}$, there exists an interval $(r_p, \bar{r}_p) \subset (p, 1)$ where (p, t) is replica breaking if $t \in (r_p, \bar{r}_p)$, and (p, t) is replica symmetric for the other values of t . Since $\phi_p(t) < I_p(t)$ if $t \in (r_p, \bar{r}_p)$ and $\phi_p(t) = I_p(t)$ for other values of t , and since $I_p(u)$ is convex, the convex minorant of $\phi_p(x)$ must touch ϕ_p at at least one $t_p \in (r_p, \bar{r}_p)$. So (p, t_p) is replica breaking and satisfies the minorant condition. ■

5 Numerical simulations using importance sampling

We implement the importance sampling schemes to show the optimality properties of the Gibbs measure tilts in practice. Although we have thus far been considering importance sampling schemes that draw i.i.d. samples from the tilted measure \mathbb{Q} , in practice it is very difficult to sample independent copies of exponential random graphs. This is because of the dependencies of the edges in the exponential random graph, unlike the situation with an Erdős-Rényi graph where the edges are independent. Thus, to implement the importance sampling scheme, we turn to a Markov chain Monte Carlo method known as the Glauber dynamics to generate samples from the exponential random graph. The Glauber dynamics refers to a Markov chain whose stationary distribution is the Gibbs measure $\mathbb{Q}_n^{h,\beta,\alpha}$. The samples \tilde{X}_k from the Glauber dynamics are used to form the importance sampling estimator \tilde{M}_K in (3.2). The variance of \tilde{M}_K clearly also depends on the correlation between the successive samples. However, in this paper, rather than focus on the effect of correlation on the variance of \tilde{M}_K , we instead investigate and compare the optimality of the importance sampling schemes, and show that importance sampling is a viable method for moderate values of n .

Glauber dynamics.

For the exponential random graph $\mathcal{G}_n^{h,\beta,\alpha}$, the Glauber dynamics proceeds as follows.

Suppose we have a graph $X = (X_{ij})_{1 \leq i < j \leq n}$. The graph \tilde{X} is generated from X via the following procedure.

1. Choose an edge X_{ij} , for some (i, j) , from X uniformly at random.
2. For the new graph \tilde{X} , fix all other edges $\tilde{X}_{i'j'} = X_{i'j'}$, for $(i', j') \neq (i, j)$.
3. Conditioned on all other edges fixed, pick

$$\tilde{X}_{ij} \sim \text{Bern}(\varphi)$$

where

$$\varphi = \frac{e^{h+(\beta/n)(L_{ij}+M_{ij})^\alpha (n^3/6)^{1-\alpha}}}{e^{h+(\beta/n)(L_{ij}+M_{ij})^\alpha (n^3/6)^{1-\alpha}} + e^{(\beta/n)M_{ij}^\alpha (n^3/6)^{1-\alpha}}}$$

and where

$$L_{ij} = \sum_{k \neq i,j} X_{ik}X_{jk}, \quad \text{and} \quad M_{ij} = \sum_{(k,l,m) \not\subseteq (i,j)} X_{kl}X_{km}X_{lm},$$

is the number of 2-stars in X with a base at the edge X_{ij} , and the number of triangles in X not involving the edge X_{ij} , respectively.

4. If conditioning on A_J is used, check if \tilde{X} is in A_J . If not, revert to X .

In step 4, a conditioning of the Gibbs measure is discussed in Section 5.2.

For the classical exponential random graph with $\alpha = 1$, the probability φ in the Glauber dynamics has a neater expression,

$$\varphi = \frac{e^{h+\beta L_{ij}/n}}{1 + e^{h+\beta L_{ij}/n}}.$$

At each MCMC step, if $X_{ij} \neq \tilde{X}_{ij}$, then $E(\tilde{X})$ differs from $E(X)$ by one edge, and $T(\tilde{X})$ differs from $T(X)$ by nL_{ij} triangles. The stationary distribution of the Glauber dynamics is the Gibbs measure $\mathbb{Q}_n^{h,\beta,\alpha}$ that defines the exponential random graph $\mathcal{G}_n^{h,\beta,\alpha}$. Regarding the mixing time of the Glauber dynamics, [1] showed for the case $\alpha = 1$ that if the variational form for the free energy of the Gibbs measure $\mathbb{Q}_n^{h,\beta,1}$,

$$\sup_{0 \leq u \leq 1} \left[\frac{h}{2}u + \frac{\beta}{6}u^3 - I(u) \right], \tag{5.1}$$

has a unique local maximum, then the mixing time of the Glauber dynamics is $\mathcal{O}(n^2 \log n)$; otherwise, the variational form has multiple local maxima, and the mixing time is $\mathcal{O}(e^n)$.

Clearly, the importance sampling tilt must be chosen so that the mixing time of the Glauber dynamics is $\mathcal{O}(n^2 \log n)$.

5.1 Example 1

The importance sampling scheme was performed for $p = 0.35, t = 0.4$, in the replica symmetry phase. We use the Glauber dynamics to draw samples from the edge tilt with parameters $h = h_t, \beta = 0$, as well as from the triangle tilt with parameters $h = h_p, \alpha = 1$ and $\beta = \frac{h_t - h_p}{t^2}$ as in (4.1). The mixing time for both tilts is $\mathcal{O}(n^2 \log n)$. In addition to the edge and triangle tilts, we also consider a family of “hybrid” tilts with parameters $h = h_q$ for $q \geq p$, and $\alpha = 1$ and

$$\beta = \beta_q = \frac{h_t - h_q}{t^2}. \tag{5.2}$$

With these parameters, the variational form for the free energy of the corresponding Gibbs measure is uniquely maximized at t . Thus, the hybrid tilt satisfies the necessary condition in Proposition 3.1 for optimality (i.e., (3.3) does not hold). By Theorem 2.4, the corresponding exponential random graph $\mathcal{G}_n^{h_q,\beta_q,1}$ is indistinguishable from the Erdős-Rényi graph $\mathcal{G}_{n,t}$ and has a mean triangle density of t^3 , in the sense of (2.22).

In the simulations, we used the hybrid tilts with $h = h_q$, for $q = 0.35, 0.36, \dots, 0.4$, and β_q satisfying (5.2). With this notation, in fact, $q = p = 0.35$ corresponds to the triangle tilt while $q = t = 0.4$ corresponds to the edge tilt. Table 1 verifies the accuracy of the importance sampling estimates for $\mu_n := \mathbb{P}(\mathcal{G}_{n,p} \in \mathcal{W}_t)$ using the tilts $\mathbb{Q}_n^{h_q,\beta_q,1}$. Also shown is the estimate for the log probability, $\frac{1}{n^2} \log \mathbb{P}(\mathcal{G}_{n,p} \in \mathcal{W}_t)$, which can be seen to approach the LDP rate

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(\mathcal{G}_{n,p} \in \mathcal{W}_t) = -I_p(t) \approx -0.002694.$$

as n is increased. Table 2 shows the estimated values of the variance of the estimator, $\text{Var}_{\mathbb{Q}_n}(\hat{q}_n)$, where $\hat{q}_n = \mathbf{1}_{\mathcal{W}_t} \frac{d\mathbb{P}_{n,p}}{d\mathbb{Q}_n^{h,\beta}}$, as well as the log second moment $\frac{1}{n^2} \log \mathbb{E}_{\mathbb{Q}_n}[\hat{q}_n^2]$. The variance of the estimator for all the hybrid and edge tilts appear to be comparable to the optimal triangle tilt, and the log second moment likewise appears to converge towards $-2I_p(t) \approx -0.0053869$. Notice that the parameters $p = 0.35, t = 0.4$ do not satisfy the hypothesis of Proposition 3.5, so the assertion of non-optimality of the edge tilt may not apply in this case. Regardless, non-optimality is not apparent for mid-sized graphs up to $n = 96$.

For $n = 16, 32, 64$, the number of MCMC samples used was $5 \times 10^4 n^2 \log n$, while for $n = 96$, the number of MCMC samples used was $10^5 n^2 \log n$.

Both the random graphs corresponding to the triangle or edge tilts are expected by (2.22), (2.23) to have triangle density of t^3 and edge density of t , on average. However, there is a difference between the way that the triangle and edge tilts produce events in $\{\mathcal{T}(f) \geq t^3\}$, which is that the edge tilt tends to produce more successful samples in $\{\mathcal{T}(f) \geq t^3\}$ with higher edge density, compared to the triangle tilt. (See Figure 2.) This

Importance Sampling for rare events in Erdős-Rényi graphs

q	0.35	0.36	0.37	0.38	0.39	0.4
n						
16	0.12475 (-0.008131)	0.1247 (-0.008132)	0.12521 (-0.008116)	0.12425 (-0.008146)	0.12441 (-0.008141)	0.12435 (-0.008143)
32	0.01107 (-0.004398)	0.011056 (-0.004399)	0.011116 (-0.004394)	0.010941 (-0.004409)	0.010972 (-0.004407)	0.010729 (-0.004429)
64	2.1919e-06 (-0.003181)	2.0283e-06 (-0.003200)	2.6073e-06 (-0.003139)	5.3287e-07 (-0.003527)	1.3822e-06 (-0.003294)	3.5772e-06 (-0.003062)
96	1.1036e-11 (-0.002738)	1.6868e-11 (-0.002692)	2.0805e-11 (-0.002669)	4.4039e-11 (-0.002587)	2.6124e-11 (-0.002644)	4.497e-11 (-0.002585)

Table 1: Comparison of the estimates for the probability μ_n (top number) for varying tilts with parameters $(h_q, \beta, 1)$, where β is defined in (5.2). Also shown is the log probability $\frac{1}{n^2} \log \mathbb{P}(\mathcal{G}_{n,p} \in \mathcal{W}_t)$ (lower number).

$n \setminus q$	0.35	0.36	0.37	0.38	0.39	0.4
16	0.030839 (-0.01199)	0.03007 (-0.01206)	0.030173 (-0.01204)	0.030553 (-0.01203)	0.031902 (-0.01191)	0.034105 (-0.01173)
32	0.00039386 (-0.007391)	0.00038614 (-0.007407)	0.00040055 (-0.007377)	0.00042058 (-0.007347)	0.00047462 (-0.007253)	0.00052598 (-0.00718)
64	2.9982e-11 (-0.005879)	2.5716e-11 (-0.005917)	4.6804e-11 (-0.005774)	2.3144e-12 (-0.006513)	1.9783e-11 (-0.005995)	1.8035e-10 (-0.005461)
96	1.157e-21 (-0.005220)	2.7721e-21 (-0.005125)	4.9044e-21 (-0.005065)	2.8661e-20 (-0.004876)	1.4562e-20 (-0.004951)	6.8628e-20 (-0.004785)

Table 2: Comparison of the estimates for the variance $Var_{\mathbb{Q}}(\hat{q}_n)$ (top number) for varying tilts with parameters $(h_q, \beta, 1)$, where β is defined in (5.2). Also shown is the log second moment $\frac{1}{n^2} \log \mathbb{E}_{\mathbb{Q}}[\hat{q}_n^2]$ (lower number).

Importance Sampling for rare events in Erdős-Rényi graphs




Figure 2: Histogram of the number of edges in the samples obtained using the importance sampling scheme based on the triangle tilt (solid red line) and edge tilts (dashed blue line), conditioned on the rare event $\{T(X) \geq \binom{n}{3}t^3\}$. The dotted green line in the top left panel shows the histogram for direct Monte Carlo sampling. The vertical line indicates the expected number of edges of the graph $\mathcal{G}_{n,p}$ conditioned on the rare event.

is attributable to the fact that the edge tilt penalizes successful samples that contain the desired triangle density but with lower than expected edge density.

5.2 Example 2: Using $\alpha \neq 1$ or conditioned Gibbs measures

The importance sampling scheme was next performed for $p = 0.2, t = 0.3$, in the replica symmetric phase. We again use the Glauber dynamics to draw samples from the edge tilt with parameters $h = h_t, \beta = 0$; the mixing time here is $\mathcal{O}(n^2 \log n)$. In contrast to the previous example, for the triangle tilt with $\alpha = 1$, the variational form (5.1) has two local maxima, resulting in a mixing time of $\mathcal{O}(e^n)$. Instead, we will use a triangle tilt with $\alpha = 2/3$. Thanks to the fact that (p, t) is in the replica symmetric phase and ϕ_p is differentiable at t , Proposition 4.3 implies that for $\alpha = 2/3$, we choose

$$\beta = \frac{h_t - h_p}{(2/3)t}.$$

The simulation results for the importance sampling scheme using the triangle tilt with $\alpha = 2/3$ is shown in Tables 3, 4 and compared with the results for the edge tilt. The simulation using direct Monte Carlo sampling is also shown for $n = 32$. We see that the triangle tilt with $\alpha = 2/3$ outperforms both the edge tilt and direct Monte Carlo simulation. Notice that the parameters $p = 0.2, t = 0.3$ do not satisfy the hypothesis of Proposition 3.5, but the edge tilt already appears to be non-optimal for mid-sized graphs up to $n = 64$.

Alternatively, we also consider a modification to the triangle tilt with $\alpha = 1$ and $\beta_p = \frac{h_t - h_p}{t^2}$ as in (4.1). This modification draws samples from the Gibbs measure $\mathbb{Q}_n^{h_p, \beta_p, 1}$ conditioned on the event that the edge and triangle densities not exceed a given threshold

n	Triangle tilt $\alpha = 2/3$	Conditioned triangle tilt	Edge tilt	Monte Carlo
16	0.0064 (-0.0197)	0.006474 (-0.0197)	0.006285 (-0.0198)	
32	4.3148e-7 (-0.0143)	3.5488e-7 (-0.0145)	3.3878e-7 (-0.0145)	3.7758e-7 (-0.0144)
48	1.3976e-13 (-0.0128)	1.1418e-14 (-0.0139)	1.2039e-12 (-0.0119)	— —
64	6.1882e-21 (-0.0136)	2.9076e-23 (-0.0127)	1.8316e-19 (-0.0105)	— —

Table 3: Estimates for the probability μ_n . In parenthesis is the estimator for the log probability $\frac{1}{n^2} \log \mu_n$.

n	Triangle tilt $\alpha = 2/3$	Conditioned triangle tilt	Edge tilt	Monte Carlo
16	1.5059e-4 (-0.0334)	2.4166e-4 (-0.0319)	1.0391e-3 (-0.0267)	
32	1.5222e-12 (-0.0265)	1.9083e-12 (-0.0263)	6.7116e-11 (-0.0229)	3.7758e-7 (-0.0144)
48	2.6058e-25 (-0.0245)	3.268e-27 (-0.0265)	2.4737e-20 (-0.0196)	— —
64	7.1703e-40 (-0.0220)	2.8806e-44 (-0.0245)	1.2806e-33 (-0.0185)	— —

Table 4: Estimates for the variance $Var_{Q_n}(\hat{q}_n)$. In parenthesis is the estimate for the log second moment, $\frac{1}{n^2} \log \mathbb{E}_{Q_n}[\hat{q}_n^2]$.

r . To be specific, let

$$A_r = \{f \in \mathcal{W} : \mathcal{T}(f) \leq r^3 \text{ and } \mathcal{E}(f) \leq r\} \tag{5.3}$$

for some $r > t$, and let the conditioned triangle tilt be defined by the Gibbs measure conditioned on A_r ,

$$\tilde{Q}_{n,A_r}^{\beta_p, \beta_p, 1}(X) \propto \begin{cases} e^{n^2(\frac{1}{2}\mathcal{E}(X) + \frac{\beta}{8}\mathcal{T}(X))}, & \text{if } X \in A_r \\ 0 & \text{if } X \notin A_r \end{cases} \tag{5.4}$$

In the numerical simulations, the threshold is chosen to be $r \approx 0.4272 > t$, which is a local minimum of the variational form (5.1). The motivation for this choice of threshold r is discussed in the Appendix. The results for the conditioned triangle tilt are shown in Tables 3, 4, which indicate that both triangle tilts perform comparably and both outperform the edge tilt.

Acknowledgments. We thanks a referee for providing many helpful comments which significantly improved the paper. SB was partially supported by was partially supported by NSF grants DMS-1105581, DMS-1310002 and SES grant 1357622. JH was partially supported by NSF grants DMS-1016441 and DMS-1512945. JN was partially supported by NSF grant DMS-1007572. This work was funded in part through the 2011-2012 SAMSI Program on Uncertainty Quantification, in which each of the authors participated and which provided postdoctoral funding for CL.

A Characterizing the phase diagrams

We present in this appendix section a framework to define subregimes of the (p, t) phase space, which extends the set up from [17].




Figure 3: Phase diagram for $t \geq p$. \mathcal{S}_1 is the dark gray region to the right of the solid curve. $\mathcal{S}_{2/3}$, the replica symmetric phase, covers both the light and dark gray regions. The diagonal red dotted curve shows $G(t) = 0$, where G is the function in (3.13). For parameters above this curve, the edge tilt does not give an optimal importance sampling estimator.

Recall that (p, t) satisfies the *minorant condition with parameter α* if the point $(t^{3\alpha}, \phi_p(t))$ lies on the convex minorant of the function $x \mapsto \phi_p(x^{1/3\alpha})$. Using the minorant condition and Lemma 4.4, we define a parameterized family of subregimes of the (p, t) -phase space.

Definition A.1. Let $\alpha \in [0, 1]$. We define the regime \mathcal{S}_α to be the set of parameters (p, t) for which the minorant condition holds with α . ■

[17] has characterized the replica symmetric phase in terms of \mathcal{S}_α .

Lemma A.2. [17, Theorems 1.1 and 1.3] $\mathcal{S}_{2/3}$ is exactly the replica symmetric phase.

Figure 3 shows the replica symmetry phase $\mathcal{S}_{2/3}$, as well as the region \mathcal{S}_1 . Using the Definition A.1, we can rephrase Prop 4.3 and Lemma 4.2 as follows.

Corollary A.3. If $(p, t) \in \mathcal{S}_\alpha$, then there exists a triangle tilt that produces an optimal scheme.

Moreover, the regime $\bigcup_{\alpha > 0} \mathcal{S}_\alpha$ where an optimal triangle tilt exists is strictly larger than the replica symmetric phase, and contains a nontrivial subset of the replica breaking phase.

Finally, we show in the next lemma that the union, $\bigcup_{\alpha > 0} \mathcal{S}_\alpha$, is an increasing union as $\alpha \rightarrow 0$. A consequence of this lemma is that if $(p, t) \in \mathcal{S}_\alpha$ for some $\alpha \in [0, 1]$, then for any $\alpha' \in [0, \alpha]$, the triangle tilt with parameters (h_p, β', α') also produces an optimal scheme, where β' is appropriately chosen depending on α' .

Lemma A.4. Let \mathcal{S}_α be defined in Definition A.1. Then $\mathcal{S}_{\alpha'} \subset \mathcal{S}_\alpha$ for $0 < \alpha < \alpha'$.

Proof. Denote $\phi_p^\alpha(x) = \phi_p(x^{1/3\alpha})$ and let $\hat{\phi}_p^\alpha(x)$ be the convex minorant of $\phi_p^\alpha(x)$. Then

$$\phi_p^{\alpha'}(x) = \phi_p(x^{1/3\alpha'}) = \phi_p((x^{\alpha/\alpha'})^{1/3\alpha}) = \phi_p^\alpha(x^{\alpha/\alpha'}).$$

Define $\eta(x) = \hat{\phi}_p^{\alpha'}(x^{\alpha'/\alpha})$. Let K be the set where $\eta(x) = \phi_p^{\alpha'}(x^{\alpha'/\alpha})$ for $x \in K$. Then

$$\eta(x) \leq \phi_p^{\alpha'}(x^{\alpha'/\alpha}) = \phi_p^\alpha(x)$$

with equality occurring iff $x \in K$. (The interpretation of K is that $t^{3\alpha} \in K$ if and only if (p, t) satisfies the minorant condition with α' .) Since $\frac{\alpha'}{\alpha} > 1$, the function $\eta(x)$ is convex and is less than $\phi_p^\alpha(x)$, hence it must be less than the convex minorant, $\eta(x) \leq \hat{\phi}_p^\alpha(x)$. For $x \in K$,

$$\phi_p^\alpha(x) = \eta(x) \leq \hat{\phi}_p^\alpha(x) \leq \phi_p^\alpha(x)$$

so $(x, \phi_p^\alpha(x))$ lies on the convex minorant $\hat{\phi}_p^\alpha(x)$ for all $x \in K$. Hence, if (p, t) satisfying the minorant condition with α' , then $t^{3\alpha} \in K$ and $(t^{3\alpha}, \phi_p(t))$ lies on the convex minorant $\hat{\phi}_p^\alpha(x)$, implying that (p, t) satisfies the minorant condition with α .

Now let (p, t) satisfy the minorant condition with α' , and suppose that $\frac{\beta'}{6}$ is a subdifferential of $\hat{\phi}_p^{\alpha'}(x)$ at the point $t^{3\alpha'}$ such that $\sup[\frac{\beta'}{6}u^{3\alpha} - \phi_p(u)]$ is uniquely maximized at t . According to the arguments in the proof of Lemma 4.4, this means that the subtangent line

$$\ell_{\alpha'}(x) := \frac{\beta'}{6}(x - t^{3\alpha'}) - \phi_p(t)$$

lies below $\phi_p^{\alpha'}(x)$ and touches it at exactly one point $t^{3\alpha'}$. Let $\nu(x) = \ell_{\alpha'}(x^{\alpha'/\alpha})$. We have that $\nu(t^{3\alpha}) = \phi_p^{\alpha'}(t^{3\alpha'}) = \phi_p^\alpha(t^{3\alpha})$ and $\nu'(t^{3\alpha}) = \frac{\beta'}{6} \frac{\alpha'}{\alpha} t^{3(\alpha'-\alpha)}$. Since $\frac{\alpha'}{\alpha} > 1$, $\nu(x)$ is convex, and the line

$$\ell_\alpha(x) := \frac{\beta}{6}(x - t^{3\alpha}) - \phi_p(t),$$

where $\frac{\beta}{6} = \nu'(t^{3\alpha})$, is tangent to $\nu(x)$ at the point $t^{3\alpha}$ and lies below $\nu(x)$. For $x \neq t^{3\alpha}$,

$$\nu(x) = \ell_{\alpha'}(x^{\alpha'/\alpha}) < \phi_p^{\alpha'}(x^{\alpha'/\alpha}) = \phi_p^\alpha(x),$$

so $\ell_\alpha(x)$ lies below $\phi_p^\alpha(x)$ and touches it at exactly one point $t^{3\alpha}$. Moreover, since $\nu(x)$ is a convex function less than $\phi_p^\alpha(x)$, we have $\hat{\phi}_p^\alpha(x) \geq \nu(x) \geq \ell_\alpha(x)$. So, $\frac{\beta}{6}$ is a subdifferential of $\hat{\phi}_p^\alpha(x)$ and $\sup[\frac{\beta}{6}u^{3\alpha} - \phi_p(u)]$ is uniquely maximized at t . The proof is complete. ■

Remark A.5. In Prop 3.5, the critical value, $\tilde{p} = \frac{e^{-1/2}}{1+e^{-1/2}} \approx 0.3775$, corresponds to $h_{\tilde{p}} = -1/2$. We see from Figure 3 that the conditions of the proposition are attained when if $p < \tilde{p}$ and (p, t) is in the region above the red dotted line intersected with the replica symmetric phase. In this region, we have $G(t) < 0$, where G is defined at (3.13). The edge tilt $\mathbb{Q}_n^{h_t, 0}$ does not produce an optimal estimator for the parameters in this region.

B Sampling from a conditioned Gibbs measure

For exponential random graphs with $\alpha = 1$, the Glauber dynamics is known to have an exponential mixing time $\mathcal{O}(e^n)$ when the variational form (5.1) has multiple local maxima [1]. When considering a triangle tilt whose variational form has multiple local maxima, the slow mixing is one reason to preclude its feasibility as an importance sampling tilt. Another reason to avoid this tilt is because the global maximum of the variational form may not occur at t , even though it has a local maximum at t by definition. Due to the second reason, such a tilt may produce a large number of samples with an over- or under-abundance of triangles, where the triangle density is determined by the global maximum of the variational form, rather than by the local minimum at t . This leads to a poor estimator that is not optimal and has large variance.

We propose to circumvent these problems by modifying the triangle tilt so that the sampled graphs are restricted to the subregion of the state space that has just the ‘right’ number of triangles.

Conditioned Gibbs measure

Given a set $A \subset \mathcal{W}$, the exponential random graph conditioned on A , denoted $\mathcal{G}_{n,A}^{h,\beta,\alpha}$ has the conditional Gibbs measure

$$\tilde{\mathbb{Q}}_{n,A}^{h,\beta,\alpha}(X) \propto \begin{cases} e^{n^2\mathcal{H}(X)}, & \text{if } X \in A \\ 0 & \text{if } X \notin A \end{cases}$$

where the Hamiltonian $\mathcal{H}(X)$ is defined in (2.9). The asymptotic behaviour of the free energy of the conditional Gibbs measure,

$$\tilde{\psi}_{n,A}^{h,\beta,\alpha} = \frac{1}{n^2} \log \sum_{X \in A} e^{n^2\mathcal{H}(X)},$$

is described in the following proposition, which follows from a direct modification of [7, Theorems 3.1, 3.2].

Proposition B.1. *For any bounded continuous mapping $\mathcal{H} : \mathcal{W} \mapsto \mathbb{R}$, and any closed subset $A \subset \mathcal{W}$, let $\tilde{\psi}_{n,A} = \tilde{\psi}_{n,A}^{h,\beta,\alpha}$ as above. Then*

$$\lim_{n \rightarrow \infty} \tilde{\psi}_{n,A} = \sup_{f \in A} [\mathcal{H}(f) - \mathcal{I}(f)]. \tag{B.1}$$

Moreover, if the variational form is maximized on the set $\tilde{\mathcal{F}} \subset A$, then the corresponding conditioned exponential random graph is asymptotically indistinguishable from $\tilde{\mathcal{F}}$.

As a consequence of Proposition B.1, an argument akin to the proof of Theorem 2.4 implies that for (p, t) in the replica symmetric phase, the conditioned Gibbs measure (5.4) conditioned on A_r has free energy given by

$$\lim_{n \rightarrow \infty} \tilde{\psi}_{n,A_r} = \sup_{0 \leq u \leq r} \left[\frac{h_p}{2} u + \frac{\beta}{6} u^3 - I(u) \right]. \tag{B.2}$$

By choosing r so that the supremum is attained at t , the corresponding exponential random graph conditioned in A_r is asymptotically indistinguishable from the Erdős-Rényi graph $\mathcal{G}_{n,t}$. Thus, the necessary condition for asymptotic optimality of the importance sampling estimator is satisfied.

Importance sampling using the conditioned Gibbs measure

The importance sampling scheme based on the conditioned Gibbs measure $\tilde{\mathbb{Q}}_{n,A}^{h,\beta,\alpha}$ gives the estimator

$$\hat{\nu}_n = \frac{1}{K} \sum_{k=1}^K \mathbf{1}_{\mathcal{W}_t}(\tilde{X}_k) \frac{d\tilde{\mathbb{P}}_{n,p,A_r}}{d\tilde{\mathbb{Q}}_{n,A_r}^{h,\beta,\alpha}}(\tilde{X}_k), \quad \text{where } \tilde{X}_k \sim \text{i.i.d. } \tilde{\mathbb{Q}}_{n,A_r}^{h,\beta,\alpha} \tag{B.3}$$

where $\tilde{\mathbb{P}}_{n,p,A_r}$ is the measure of the Erdős-Rényi graph conditioned on A_r . Note that $\hat{\nu}_n$ is an unbiased estimator for $\nu_n = \tilde{\mathbb{P}}_{n,p,A_r}(\mathcal{W}_t)$, but it is a biased estimator for $\mu_n = \mathbb{P}_{n,p}(\mathcal{W}_t)$. The bias can be corrected by

$$\hat{\mu}_n = \hat{\nu}_n \cdot \mathbb{P}_{n,p}(A_r) + \mathbb{P}_{n,p}(\mathcal{W}_t \cap A_r^c),$$

but the two probabilities on the RHS are not be easily computable or estimated. Nonetheless, the choice of the set A_r ensures the bias is small and vanishes exponentially faster than the small probability μ_n . In fact, standard computations give that $\mathbb{P}_{n,p}(A_r) \rightarrow 1$ as $n \rightarrow \infty$, and

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}_{n,p}(\mathcal{W}_t \cap A_r^c) \leq - \inf_{f \in \mathcal{W}_t} [I_p(f)].$$

The asymptotic optimality of the importance sampling scheme is stated in the following result.

Corollary B.2. *Given (p, t) in the replica symmetric phase, consider the conditioned triangle tilt defined by the Gibbs measure $\mathbb{Q}_{n, A_r}^{h_p, \beta_p, 1}$ in (5.4), conditioned on A_r in (5.3) with $p < t < r$. The importance sampling scheme based on this conditioned triangle tilt is asymptotically optimal.*

Choosing the set A_r .

We motivate the choice of the set A_r in the Example from Section 5.2, with $p = 0.2, t = 0.3$ in the replica symmetric phase. We had noted that for the triangle tilt with Gibbs measure $\mathbb{Q}_n^{h_p, \beta_p, 1}$, the variational form $V(u) = \frac{h_p}{2}u + \frac{\beta_p}{6}u^3 - I(u)$ has multiple local maxima. This is illustrated in Figure 4 (inset), where $t = 0.3$ is a local maximum but not a global maximum, whereas $u^* \approx 0.989$ is the global maximum. Without conditioning, the exponential graph $\mathcal{G}_n^{h_p, \beta_p, 1}$ has a mean triangle density of $(u^*)^3$, much greater than the desired triangle density of t^3 ; moreover, successive samples in the Glauber dynamics take exponentially long time to move from the region with a high triangle density $(u^*)^3$ to the region with a lower triangle density t^3 . The effect of conditioning on the set A_r is to cap the triangle density at r^3 , and ensure faster mixing of the Glauber dynamics. Thus, one convenient choice of r is to take $r \approx 0.4272$ to be the local minimum of $V(u)$ which separates the two local maxima. Then, t is the unique global maximum on the interval $[0, r]$ and the conditioned Gibbs measure has a mean triangle density of t^3 .

Conditioning the Gibbs measure leads to a significant reduction in the asymptotic log second moment of the importance sampling estimator. This reduction is best illustrated by considering, besides the triangle tilt itself, the family of Gibbs measures with $h = h_p, \alpha = 1$ and varying $\beta > 0$. Figure 4 illustrates that as β is increased from 0 up to a transition point $\beta \approx 4.76$, the variational form $V(u; \beta) = \frac{h_p}{2}u + \frac{\beta}{6}u^{3\alpha} - I(u)$ has a global maximum within the range $[0.2, 0.3]$. For $\beta > 4.76$, the global maximum jumps up into the range $[0.9, 1]$, so that the exponential random graph transitions from a regime of low edge density to one of high edge density. Near to the transition point, there is a range of β for which $V(u; \beta)$ has two local maxima and one local minimum. Observe from the figure inset that the triangle tilt with $\beta = \beta_p$ lies in this range. It is for this range of β that applying the conditioned Gibbs measure will lead to a reduction in the asymptotic log second moment. Figure 5 shows the asymptotic log second moment, both with and without conditioning of the Gibbs measure. For each β , the threshold r is chosen as the local minimum of $V(u; \beta)$. We see that conditioning the Gibbs measure significantly reduces the asymptotic log second moment, and the conditioned triangle tilt is asymptotically optimal. This is corroborated by the numerical simulations presented in Section 5.2. In contrast, when no conditioning is performed, the IS estimator exhibits a sharp decline in performance when β is increased beyond the transition point at $\beta \approx 4.76$.

C Auxiliary lemmas and proofs

We present a lemma on the asymptotic indistinguishability of an exponential random graph from a minimal set \mathcal{F}^* , as well as the proof of Proposition 2.6.

Lemma C.1. *[(i)]*

1. *Given (p, t) , let \mathcal{F}^* be the set of functions that minimize the LDP rate function, $\inf_{f \in \mathcal{W}_t} [\mathcal{I}_p(f)]$ in (2.5). Then \mathcal{F}^* is the minimal set that the Erdős-Rényi graph $\mathcal{G}_{n,p}$ conditioned on $\{\mathcal{T}(f) \geq t^3\}$ is asymptotically indistinguishable from.*
2. *Given (h, β, α) , let \mathcal{F}^* be the set of functions that maximize $\sup_{f \in \mathcal{W}} [\mathcal{H}(f) - \mathcal{I}(f)]$. Then \mathcal{F}^* is the minimal set that the exponential random graph $\mathcal{G}_n^{h, \beta, \alpha}$ is asymptoti-*




Figure 4: The phase curve denotes the values of the stationary points of the variational form $V(u) = \frac{h}{2}u + \frac{\beta}{6}u^{3\alpha} - I(u)$, as β varies, and given $\alpha = 1$, $h_p = \log \frac{p}{1-p}$, $p = 0.2$. The red solid line denotes when the stationary point is a global maximum of $V(u)$; the red dotted line denotes the local maximum; the blue dashed line denotes the local minimum. At the phase transition point at $\beta \approx 4.76$, the maximum of the variational form jumps from $u^* \approx 0.253$ to $u^* \approx 0.947$. The inset shows the function $V(u)$ for $\beta = \beta^* \approx 5.99$ attaining a local maximum at $t = 0.3$ and global maximum at $u^* \approx 0.989$.




Figure 5: A plot of the asymptotic second moment, $\lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}_{\mathbb{Q}}[\hat{q}_{n,A_r}^2]$, of the importance sampling estimator based on the conditioned Gibbs tilt for fixed $h = h_p$ and varying β . The insert is a zoom-in to show that the smallest variance is attained at $\beta = \beta^*$. The dotted line shows the rapid deterioration of the asymptotic second moment of the estimator without the use of conditioning. Parameters used are $p = 0.2$ and $t = 0.3$.

cally indistinguishable from.

Proof. The proofs of asymptotic indistinguishability of \mathcal{F}^* was shown in [8, Theorem 3.1] for (i) and [7, Theorem 3.22] for (ii). The proofs naturally extend to give the minimality of \mathcal{F}^* , and we state them here for the record.

Observe that for any random graph \mathcal{G}_n that is asymptotically indistinguishable from a set \mathcal{F}^* , to show that \mathcal{F}^* is minimal, it suffices to show that, for any relatively open non-empty subset $\mathcal{F}_0 \subset \mathcal{F}^*$ such that $\mathcal{F}^* \setminus \mathcal{F}_0$ is non-empty, there exists $\epsilon > 0$ such that

$$\liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(\delta_{\square}(\mathcal{G}_n, \mathcal{F}^* \setminus \mathcal{F}_0) > \epsilon) = 0. \tag{C.1}$$

Let $\mathcal{F}_0 \subset \mathcal{F}^*$ be any relatively open non-empty subset, with $\mathcal{F}^* \setminus \mathcal{F}_0$ non-empty. Denote, for $\epsilon > 0$,

$$\mathcal{F}_{\epsilon} = \{f \in \mathcal{W} \mid \delta_{\square}(f, \mathcal{F}^* \setminus \mathcal{F}_0) > \epsilon\}.$$

(i) Since \mathcal{F}_0 is relatively open in \mathcal{F}^* , $\delta_{\square}(f, \mathcal{F}^* \setminus \mathcal{F}_0) > 0$ for any $f \in \mathcal{F}_0$. So, there exists an $\epsilon > 0$ sufficiently small such that $(\mathcal{F}_{\epsilon} \cap \mathcal{W}_t)^{\circ}$ contains at least one element of \mathcal{F}_0 . (A° denotes the interior of A .) It follows that

$$\inf_{f \in (\mathcal{F}_{\epsilon} \cap \mathcal{W}_t)^{\circ}} [\mathcal{I}_p(f)] = \inf_{f \in \mathcal{W}_t} [\mathcal{I}_p(f)].$$

Since

$$\mathbb{P}(\mathcal{G}_{n,p} \in \mathcal{F}_{\epsilon} \mid \mathcal{G}_{n,p} \in \mathcal{W}_t) = \frac{\mathbb{P}(\mathcal{G}_{n,p} \in \mathcal{F}_{\epsilon} \cap \mathcal{W}_t)}{\mathbb{P}(\mathcal{G}_{n,p} \in \mathcal{W}_t)},$$

from the large deviation principle in [8, Theorem 2.3] implies that

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(\mathcal{G}_{n,p} \in \mathcal{F}_{\epsilon} \mid \mathcal{G}_{n,p} \in \mathcal{W}_t) \\ &= \liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(\mathcal{G}_{n,p} \in \mathcal{F}_{\epsilon} \cap \mathcal{W}_t) - \frac{1}{n^2} \log \mathbb{P}(\mathcal{G}_{n,p} \in \mathcal{W}_t) \\ &\geq - \inf_{f \in (\mathcal{F}_{\epsilon} \cap \mathcal{W}_t)^{\circ}} [\mathcal{I}_p(f)] + \inf_{f \in \mathcal{W}_t} [\mathcal{I}_p(f)] \\ &= 0. \end{aligned}$$

(ii) Since \mathcal{F}_0 is relatively open in \mathcal{F}^* , there exists an $\epsilon > 0$ sufficiently small such that $\mathcal{F}_{\epsilon}^{\circ}$ contains at least one element of \mathcal{F}_0 , and

$$\inf_{f \in \mathcal{F}_{\epsilon}^{\circ}} [\mathcal{H}(f) - \mathcal{I}(f)] = \inf_{f \in \mathcal{W}} [\mathcal{H}(f) - \mathcal{I}(f)].$$

Since the Hamiltonian \mathcal{H} is bounded, for any $\eta > 0$, there is a finite set $A \subset \mathbb{R}$ such that the intervals $\{(a, a + \eta), a \in A\}$ cover the range of \mathcal{H} . Let $\mathcal{F}_{\epsilon}^a = \mathcal{F}_{\epsilon} \cap \mathcal{H}^{-1}([a, a + \eta])$, and let $\mathcal{F}_{\epsilon}^{a,n} = \mathcal{F}_{\epsilon}^a \cap \Omega_n$ be the functions corresponding to a simple finite graph. Then

$$\mathbb{P}(\mathcal{G}_n \in \mathcal{F}_{\epsilon}) \geq \sum_{a \in A} e^{n^2(a - \psi_n)} |\mathcal{F}_{\epsilon}^{a,n}| \geq e^{-n^2 \psi_n} \sup_{a \in A} [e^{n^2 a} |\mathcal{F}_{\epsilon}^{a,n}|]$$

and

$$\frac{1}{n^2} \log \mathbb{P}(\mathcal{G}_n \in \mathcal{F}_{\epsilon}) \geq -\psi_n + \sup_{a \in A} [a - \frac{1}{n^2} \log |\mathcal{F}_{\epsilon}^{a,n}|].$$

By an observation in [7, Eqn. (3.4)], for any open set $U \subset \mathcal{W}$, and $U_n = U \cap \Omega_n$,

$$\liminf_{n \rightarrow \infty} \frac{1}{n^2} \log |U_n| \geq - \inf_{f \in U} [\mathcal{I}(f)].$$

Then, since

$$\sup_{f \in \mathcal{F}_\varepsilon^a} [\mathcal{H}(f) - \mathcal{I}(f)] \leq \sup_{f \in \mathcal{F}_\varepsilon^a} [a + \eta - \mathcal{I}(f)] = a + \eta - \inf_{f \in \mathcal{F}_\varepsilon^a} [\mathcal{I}(f)]$$

we have that

$$\begin{aligned} \liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(\mathcal{G}_n \in \mathcal{F}_\varepsilon) &\geq - \sup_{f \in \mathcal{W}} [\mathcal{H}(f) - \mathcal{I}(f)] + \sup_{a \in A} [a - \inf_{f \in (F_\varepsilon^a)^\circ} [\mathcal{I}(f)]] \\ &\geq - \sup_{f \in \mathcal{W}} [\mathcal{H}(f) - \mathcal{I}(f)] + \sup_{a \in A} \sup_{f \in (F_\varepsilon^a)^\circ} [\mathcal{H}(f) - \mathcal{I}(f)] - \eta \\ &\geq - \sup_{f \in \mathcal{W}} [\mathcal{H}(f) - \mathcal{I}(f)] + \sup_{f \in F_\varepsilon^\circ} [\mathcal{H}(f) - \mathcal{I}(f)] - \eta \\ &= 0. \end{aligned}$$

The proof is complete. ■

Proof of Proposition 2.6.

Proof. Let $\epsilon_1 > 0$ be arbitrary. As in Theorem 2.4, let $\mathcal{F}_{v^*}^*$ be the set of minimizers of $\inf_{f \in \partial \mathcal{W}_{v^*}} [\mathcal{I}_q(f)]$.

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}_n} |\mathcal{T}(X) - (v^*)^3| &= \int_{\{\delta_\square(X, \mathcal{F}_{v^*}^*) > \epsilon_1\}} |\mathcal{T}(X) - (v^*)^3| d\mathbb{Q}_n(X) + \int_{\{\delta_\square(X, \mathcal{F}_{v^*}^*) \leq \epsilon_1\}} |\mathcal{T}(X) - (v^*)^3| d\mathbb{Q}_n(X) \\ &= (I) + (II) \end{aligned}$$

(We have dropped the superscripts, $\mathbb{Q}_n = \mathbb{Q}_n^{h_q, \beta, \alpha}$.) We estimate the two terms. To estimate (I), by [7, Theorem 4.2], there exists $C, \epsilon_2 > 0$ such that for sufficiently large n

$$\mathbb{Q}_n(\delta_\square(X, \mathcal{F}_{v^*}^*) > \epsilon_1) \leq C_2 e^{-n^2 \epsilon_2}.$$

Since $|\mathcal{T}(X) - (v^*)^3| \leq 1$,

$$(I) \leq \mathbb{Q}_n(\delta_\square(X, \mathcal{F}_{v^*}^*) > \epsilon_1) \leq C_2 e^{-n^2 \epsilon_2}.$$

To estimate (II), for any $X \in \{\delta_\square(X, \mathcal{F}_{v^*}^*) \leq \epsilon_1\}$, let the function $f_X^* \in \mathcal{F}_{v^*}^*$ be such that $\delta_\square(X, f_X^*) \leq \epsilon_1$. Note that $\mathcal{T}(f_X^*) = (v^*)^3$ by definition. By Lipschitz continuity of the mapping $f \mapsto \mathcal{T}(f)$ under the cut distance metric δ_\square [3, Theorem 3.7],

$$|\mathcal{T}(X) - (v^*)^3| = |\mathcal{T}(X) - \mathcal{T}(f_X^*)| \leq C_1 \delta_\square(X, f_X^*) \leq C_1 \epsilon_1.$$

So

$$\begin{aligned} (II) &= \int_{\{\delta_\square(X, \mathcal{F}_{v^*}^*) \leq \epsilon_1\}} |\mathcal{T}(X) - (v^*)^3| d\mathbb{Q}_n(X) \\ &\leq C_1 \epsilon_1 \mathbb{Q}_n(\delta_\square(X, \mathcal{F}_{v^*}^*) \leq \epsilon_1) \\ &\leq C_1 \epsilon_1. \end{aligned}$$

Hence,

$$\lim_{n \rightarrow \infty} \mathbb{E}_{\mathbb{Q}_n} |\mathcal{T}(X) - (v^*)^3| \leq \lim_{n \rightarrow \infty} C_2 e^{-n^2 \epsilon_2} + C_1 \epsilon_1 = C_1 \epsilon_1.$$

Since ϵ_1 is arbitrary, (2.22) follows.

If (q, v^*) belongs to the replica symmetric phase, we have by Theorem 2.4 that $\mathcal{F}_{v^*}^*$ consists uniquely of the constant function $f^*(x, y) \equiv v^*$. Then since $\mathcal{E}(f^*) = v^*$, the above proof follows identically to yield that

$$\lim_{n \rightarrow \infty} \mathbb{E}_{\mathbb{Q}_n} |\mathcal{E}(X) - v^*| \leq \lim_{n \rightarrow \infty} C_2 e^{-n^2 \epsilon_2} + C_1 \epsilon_1 = C_1 \epsilon_1. \quad \blacksquare$$

References

- [1] Shankar Bhamidi, Guy Bresler, and Allan Sly, *Mixing time of exponential random graphs*, Ann. Appl. Probab. **21** (2011), no. 6, 2146–2170. MR-2895412
- [2] J. Blanchet and P. Glynn, *Efficient rare-event simulation for the maximum of heavy-tailed random walks*, The Annals of Applied Probability **18** (2008), no. 4, 1351–1378. MR-2434174
- [3] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, and K. Vesztegombi, *Convergent sequences of dense graphs. I. Subgraph frequencies, metric properties and testing*, Adv. Math. **219** (2008), no. 6, 1801–1851. MR-2455626
- [4] James Antonio Bucklew, *Introduction to rare event simulation*, Springer Series in Statistics, Springer-Verlag, New York, 2004. MR-2045385
- [5] Sourav Chatterjee, *The missing log in large deviations for triangle counts*, Random Structures Algorithms **40** (2012), no. 4, 437–451. MR-2925306
- [6] Sourav Chatterjee and Partha S. Dey, *Applications of Stein’s method for concentration inequalities*, Ann. Probab. **38** (2010), no. 6, 2443–2485. MR-2683635
- [7] Sourav Chatterjee and Persi Diaconis, *Estimating and understanding exponential random graph models*, Ann. Statist. **41** (2013), no. 5, 2428–2461. MR-3127871
- [8] Sourav Chatterjee and S. R. S. Varadhan, *The large deviation principle for the Erdős-Rényi random graph*, European J. Combin. **32** (2011), no. 7, 1000–1017. MR-2825532
- [9] B. DeMarco and J. Kahn, *Upper tails for triangles*, Random Structures Algorithms **40** (2012), no. 4, 452–459. MR-2925307
- [10] Paul Dupuis and Hui Wang, *Importance sampling, large deviations, and differential games*, Stoch. Stoch. Rep. **76** (2004), no. 6, 481–508. MR-2100018
- [11] Paul Glasserman and Yashan Wang, *Counterexamples in importance sampling for large deviations probabilities*, Ann. Appl. Probab. **7** (1997), no. 3, 731–746. MR-1459268
- [12] S. Juneja and P. Shahabuddin, *Rare event simulation techniques: An introduction and recent advances*, Simulation, Handbooks in Operations Research and Management Science (2006), 291–350.
- [13] J. H. Kim and V. H. Vu, *Divide and conquer martingales and the number of triangles in a random graph*, Random Structures Algorithms **24** (2004), no. 2, 166–174. MR-2035874
- [14] László Lovász, *Large networks and graph limits*, American Mathematical Society Colloquium Publications, vol. 60, American Mathematical Society, Providence, RI, 2012. MR-3012035
- [15] László Lovász and Balázs Szegedy, *Limits of dense graph sequences*, J. Combin. Theory Ser. B **96** (2006), no. 6, 933–957. MR-2274085
- [16] _____, *Szemerédi’s lemma for the analyst*, Geom. Funct. Anal. **17** (2007), no. 1, 252–270. MR-2306658
- [17] Eyal Lubetzky and Yufei Zhao, *On replica symmetry of large deviations in random graphs*, Random Structures Algorithms **47** (2015), no. 1, 109–146. MR-3366814
- [18] Charles Radin and Mei Yin, *Phase transitions in exponential random graphs*, Ann. Appl. Probab. **23** (2013), no. 6, 2458–2471. MR-3127941
- [19] G. Robins, P. Pattison, Y. Kalish, and D. Lusher, *An introduction to exponential random graph (p^*) models for social networks*, Social networks **29** (2007), no. 2, 173–191.
- [20] G. Robins, T. Snijders, P. Wang, M. Handcock, and P. Pattison, *Recent developments in exponential random graph (p^*) models for social networks*, Social networks **29** (2007), no. 2, 192–215.
- [21] G. Rubino and B. Tuffin, *Rare event simulation using monte carlo methods*, Wiley Online Library, 2009.
- [22] M. Yin, *A cluster expansion approach to exponential random graph models*, Journal of Statistical Mechanics: Theory and Experiment **2012** (2012), no. 05, P05004.
- [23] Mei Yin, *A cluster expansion approach to renormalization group transformations*, J. Math. Phys. **52** (2011), no. 3, 033502, 12. MR-2814863